

COMPUTATIONAL INTELLIGENCE AND ITS
APPLICATIONS SERIES

Advances in Applied Artificial Intelligence



JOHN FULCHER

VISIT...

LANZAROTE
Caliente.COM

Advances in Applied Artificial Intelligence

John Fulcher, University of Wollongong, Australia



IDEA GROUP PUBLISHING

Hershey • London • Melbourne • Singapore

Acquisitions Editor: Michelle Potter
Development Editor: Kristin Roth
Senior Managing Editor: Amanda Appicello
Managing Editor: Jennifer Neidig
Copy Editor: Susanna Svidunovich
Typesetter: Sharon Berger
Cover Design: Lisa Tosheff
Printed at: Yurchak Printing Inc.

Published in the United States of America by
Idea Group Publishing (an imprint of Idea Group Inc.)
701 E. Chocolate Avenue, Suite 200
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@idea-group.com
Web site: <http://www.idea-group.com>

and in the United Kingdom by
Idea Group Publishing (an imprint of Idea Group Inc.)
3 Henrietta Street
Covent Garden
London WC2E 8LU
Tel: 44 20 7240 0856
Fax: 44 20 7379 0609
Web site: <http://www.eurospanonline.com>

Copyright © 2006 by Idea Group Inc. All rights reserved. No part of this book may be reproduced, stored or distributed in any form or by any means, electronic or mechanical, including photocopying, without written permission from the publisher.

Product or company names used in this book are for identification purposes only. Inclusion of the names of the products or companies does not indicate a claim of ownership by IGI of the trademark or registered trademark.

Library of Congress Cataloging-in-Publication Data

Advances in applied artificial intelligence / John Fulcher, editor.
p. cm.

Summary: "This book explores artificial intelligence finding it cannot simply display the high-level behaviours of an expert but must exhibit some of the low level behaviours common to human existence"--Provided by publisher.

Includes bibliographical references and index.

ISBN 1-59140-827-X (hardcover) -- ISBN 1-59140-828-8 (softcover) -- ISBN 1-59140-829-6 (ebook)

1. Artificial intelligence. 2. Intelligent control systems. I. Fulcher, John.
Q335.A37 2006
006.3--dc22

2005032066

British Cataloguing in Publication Data

A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book is new, previously-unpublished material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

IGP Forthcoming Titles in the *Computational Intelligence and Its Applications Series*

Biometric Image Discrimination Technologies (February 2006 release)

David Zhang, Xiaoyuan Jing and Jian Yang

ISBN: 1-59140-830-X

Paperback ISBN: 1-59140-831-8

eISBN: 1-59140-832-6

Computational Economics: A Perspective from Computational Intelligence

(November 2005 release)

Shu-Heng Chen, Lakhmi Jain, and Chung-Ching Tai

ISBN: 1-59140-649-8

Paperback ISBN: 1-59140-650-1

eISBN: 1-59140-651-X

Computational Intelligence for Movement Sciences: Neural Networks, Support Vector Machines and Other Emerging Technologies (February 2006 release)

Rezaul Begg and Marimuthu Palaniswami

ISBN: 1-59140-836-9

Paperback ISBN: 1-59140-837-7

eISBN: 1-59140-838-5

An Imitation-Based Approach to Modeling Homogenous Agents Societies

(July 2006 release)

Goran Trajkovski

ISBN: 1-59140-839-3

Paperback ISBN: 1-59140-840-7

eISBN: 1-59140-841-5

It's Easy to Order! Visit www.idea-group.com!

717/533-8845 x10

Mon-Fri 8:30 am-5:00 pm (est) or fax 24 hours a day 717/533-8661



IDEA GROUP PUBLISHING

Hershey • London • Melbourne • Singapore

Excellent additions to your library!

*This book is dedicated to
Taliver John Fulcher.*

Advances in Applied Artificial Intelligence

Table of Contents

Preface viii

Chapter I

Soft Computing Paradigms and Regression Trees in Decision Support Systems 1

Cong Tran, University of South Australia, Australia

Ajith Abraham, Chung-Ang University, Korea

Lakhmi Jain, University of South Australia, Australia

Chapter II

Application of Text Mining Methodologies to Health Insurance Schedules 29

Ah Chung Tsoi, Monash University, Australia

Phuong Kim To, Tedis P/L, Australia

Markus Hagenbuchner, University of Wollongong, Australia

Chapter III

Coordinating Agent Interactions Under Open Environments 52

Quan Bai, University of Wollongong, Australia

Minjie Zhang, University of Wollongong, Australia

Chapter IV

Literacy by Way of Automatic Speech Recognition 68

Russell Gluck, University of Wollongong, Australia

John Fulcher, University of Wollongong, Australia

Chapter V

Smart Cars: The Next Frontier 120

Lars Petersson, National ICT Australia, Australia

Luke Fletcher, Australian National University, Australia

Nick Barnes, National ICT Australia, Australia

Alexander Zelinsky, CSIRO ICT Centre, Australia

Chapter VI

The Application of Swarm Intelligence to Collective Robots 157

Amanda J. C. Sharkey, University of Sheffield, UK

Noel Sharkey, University of Sheffield, UK

Chapter VII

Self-Organising Impact Sensing Networks in Robust Aerospace Vehicles 186

Mikhail Prokopenko, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Geoff Poulton, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Don Price, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Peter Wang, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Philip Valencia, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Nigel Hoschke, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Tony Farmer, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Mark Hedley, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Chris Lewis, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Andrew Scott, CSIRO Information and Communication

Technology Centre and CSIRO Industrial Physics, Australia

Chapter VIII

Knowledge Through Evolution 234

Russell Beale, University of Birmingham, UK

Andy Pryke, University of Birmingham, UK

Chapter IX

Neural Networks for the Classification of Benign and Malignant Patterns in

Digital Mammograms 251

Brijesh Verma, Central Queensland University, Australia

Rinku Panchal, Central Queensland University, Australia

Chapter X

Swarm Intelligence and the Taguchi Method for Identification of Fuzzy Models 273

Arun Khosla, National Institute of Technology, Jalandhar, India

Shakti Kumar, Haryana Engineering College, Jalandhar, India

K. K. Aggarwal, GGS Indraprastha University, Delhi, India

About the Authors 296

Index 305

Preface

Discussion on the nature of intelligence long pre-dated the development of the electronic computer, but along with that development came a renewed burst of investigation into what an artificial intelligence would be. There is still no consensus on how to define artificial intelligence: Early definitions tended to discuss the type of behaviours which we would class as intelligent, such as a mathematical theorem proving or displaying medical expertise of a high level. Certainly such tasks are signals to us that the person exhibiting such behaviours is an expert and deemed to be engaging in intelligent behaviours; however, 60 years of experience in programming computers has shown that many behaviours to which we do not ascribe intelligence actually require a great deal of skill. These behaviours tend to be ones which all normal adult humans find relatively easy, such as speech, face recognition, and everyday motion in the world. The fact that we have found it to be extremely difficult to tackle such mundane problems suggests to many scientists that an artificial intelligence cannot simply display the high-level behaviours of an expert but must, in some way, exhibit some of the low-level behaviours common to human existence.

Yet this stance does not answer the question of what constitutes an artificial intelligence but merely moves the question to what common low-level behaviours are necessary for an artificial intelligence. It seems unsatisfactory to take the stance which some do, that states that we would know one if we met one. This book takes a very pragmatic approach to the problem by tackling individual problems and seeking to use tools from the artificial intelligence community to solve these problems. The techniques that are used tend to be those which are suggested by human life, such as artificial neural networks and evolutionary algorithms. The underlying reasoning behind such technologies is that we have not created intelligences through such high-level techniques as logic programming; therefore, there must be something in the actuality of life itself which begets intelligence. For example, the study of artificial neural networks is both an engineering study in that some practitioners wish to build machines based on artificial neural networks which can solve specific problems, but it is also a study which gives us some insight into how our own intelligences are generated. Regardless of the reason given for this study, the common rationale is that there is something in the bricks and mortar of brains — the actual neurons and synapses — which is crucial to the display of intelligence. Therefore, to display intelligence, we are required to create machines which also have artificial neurons and synapses.

Similarly, the rationale behind agent programs is based on a belief that we become intelligent within our social groups. A single human raised in isolation will never be as intelligent as one who comes into daily contact with others throughout his or her developing life. Note that for this to be true, it is also required that the agent be able to learn in some way to modulate its actions and responses to those of the group. Therefore, a pre-programmed agent will not be as strong as an agent which is given the ability to dynamically change its behaviour over time. The evolutionary approach too shares this view in that the final population is not a pre-programmed solution to a problem, but rather emerges through the processes of survival-of-the fittest and their reproduction with inaccuracies.

Whether any one technology will prove to be the central one in creating artificial intelligence or whether a combination of technologies will be necessary to create an artificial intelligence is still an open question, so many scientists are experimenting with mixtures of such techniques.

In this volume, we see such questions implicitly addressed by scientists tackling specific problems which require intelligence with both individual and combinations of specific artificial intelligence techniques.

OVERVIEW OF THIS BOOK

In **Chapter I**, Tran, Abraham, and Jain investigate the use of multiple soft computing techniques such as neural networks, evolutionary algorithms, and fuzzy inference methods for creating intelligent decision support systems. Their particular emphasis is on blending these methods to provide a decision support system which is robust, can learn from the data, can handle uncertainty, and can give some response even in situations for which no prior human decisions have been made. They have carried out extensive comparative work with the various techniques on their chosen application, which is the field of tactical air combat.

In **Chapter II**, Tsoi, To, and Hagenbuchner tackle a difficult problem in text mining — automatic classification of documents using only the words in the documents. They discuss a number of rival and cooperating techniques and, in particular, give a very clear discussion on latent semantic kernels. Kernel techniques have risen to prominence recently due to the pioneering work of Vapnik. The application to text mining in developing kernels specifically for this task is one of the major achievements in this field. The comparative study on health insurance schedules makes interesting reading.

Bai and Zhang in **Chapter III** take a very strong position on what constitutes an agent: “An intelligent agent is a reactive, proactive, autonomous, and social entity”. Their chapter concentrates very strongly on the last aspect since it deals with multi-agent systems in which the relations between agents is not pre-defined nor fixed when it is learned. The problems of inter-agent communication are discussed under two headings: The first investigates how an agent may have knowledge of its world and what ontologies can be used to specify the knowledge; the second deals with agent interaction protocols and how these may be formalised. These are set in the discussion of a supply-chain formation.

Like many of the chapters in this volume, **Chapter IV** forms almost a mini-book (at 50+ pages), but Gluck and Fulcher give an extensive review of automatic speech recognition systems covering pre-processing, feature extraction, and pattern matching. The

authors give an excellent review of the main techniques currently used including hidden Markov models, linear predictive coding, dynamic time warping, and artificial neural networks with the authors' familiarity with the nuts-and-bolts of the techniques being evident in the detail with which they discuss each technique. For example, the artificial neural network section discusses not only the standard back propagation algorithm and self-organizing maps, but also recurrent neural networks and the related time-delay neural networks. However, the main topic of the chapter is the review of the draw-talk-write approach to literacy which has been ongoing research for almost a decade. Most recent work has seen this technique automated using several of the techniques discussed above. The result is a socially-useful method which is still in development but shows a great deal of potential.

Petersson, Fletcher, Barnes, and Zelinsky turn our attention to their Smart Cars project in **Chapter V**. This deals with the intricacies of Driver Assistance Systems, enhancing the driver's ability to drive rather than replacing the driver. Much of their work is with monitoring systems, but they also have strong reasoning systems which, since the work involves keeping the driver in the loop, must be intuitive and explanatory. The system involves a number of different technologies for different parts of the system: Naturally, since this is a real-world application, much of the data acquired is noisy, so statistical methods and probabilistic modelling play a big part in their system, while support vectors are used for object-classification.

Amanda and Noel Sharkey take a more technique-driven approach in **Chapter VI** when they investigate the application of swarm techniques to collective robotics. Many of the issues such as communication which arise in swarm intelligence mirror those of multi-agent systems, but one of the defining attributes of swarms is that the individual components should be extremely simple, a constraint which does not appear in multi-agent systems. The Sharkeys enumerate the main components of such a system as being composed of a group of simple agents which are autonomous, can communicate only locally, and are biologically inspired. Each of these properties is discussed in some detail in Chapter VI. Sometimes these techniques are combined with artificial neural networks to control the individual agents or genetic algorithms, for example, for developing control systems. The application to robotics gives a fascinating case-study.

In **Chapter VII**, the topic of structural health management (SHM) is introduced. This "is a new approach to monitoring and maintaining the integrity and performance of structures as they age and/or sustain damage", and Prokopenko and his co-authors are particularly interested in applying this to aerospace systems in which there are inherent difficulties, in that they are operating under extreme conditions. A multi-agent system is created to handle the various sub-tasks necessary in such a system, which is created using an interaction between top-down dissection of the tasks to be done with a bottom-up set of solutions for specific tasks. Interestingly, they consider that most of the bottom-up development should be based on self-organising principles, which means that the top-down dissection has to be very precise. Since they have a multi-agent system, communication between the agents is a priority: They create a system whereby only neighbours can communicate with one another, believing that this gives robustness to the whole system in that there are then multiple channels of communication. Their discussion of chaotic regimes and self-repair systems provides a fascinating insight into the type of system which NASA is currently investigating. This chapter places self-referentiability as a central factor in evolving multi-agent systems.

In **Chapter VIII**, Beale and Pryke make an elegant case for using computer algorithms for the tasks for which they are best suited, while retaining human input into any investigation for the tasks for which the human is best suited. In an exploratory data investigation, for example, it may one day be interesting to identify clusters in a data set, another day it may be more interesting to identify outliers, while a third day may see the item of interest shift to the manifold in which the data lies. These aspects are specific to an individual's interests and will change in time; therefore, they develop a mechanism by which the human user can determine the criterion of interest for a specific data set so that the algorithm can optimise the view of the data given to the human, taking into account this criterion. They discuss trading accuracy for understanding in that, if presenting 80% of a solution makes it more accessible to human understanding than a possible 100% solution, it may be preferable to take the 80% solution. A combination of evolutionary algorithms and a type of spring model are used to generate interesting views.

Chapter IX sees an investigation by Verma and Panchal into the use of neural networks for digital mammography. The whole process is discussed here from collection of data, early detection of suspicious areas, area extraction, feature extraction and selection, and finally the classification of patterns into 'benign' or 'malignant'. An extensive review of the literature is given, followed by a case study on some benchmark data sets. Finally the authors make a plea for more use of standard data sets, something that will meet with heartfelt agreement from other researchers who have tried to compare different methods which one finds in the literature.

In **Chapter X**, Khosla, Kumar, and Aggarwal report on the application of particle swarm optimisation and the Taguchi method to the derivation of optimal fuzzy models from the available data. The authors emphasize the importance of selecting appropriate PSO strategies and parameters for such tasks, as these impact significantly on performance. Their approach is validated by way of data from a rapid Ni-Cd battery charger.

As we see, the chapters in this volume represent a wide spectrum of work, and each is self-contained. Therefore, the reader can dip into this book in any order he/she wishes. There are also extensive references within each chapter which an interested reader may wish to pursue, so this book can be used as a central resource from which major avenues of research may be approached.

*Professor Colin Fyfe
The University of Paisley, Scotland
December, 2005*

Chapter I

Soft Computing Paradigms and Regression Trees in Decision Support Systems

Cong Tran, University of South Australia, Australia

Ajith Abraham, Chung-Ang University, Korea

Lakhmi Jain, University of South Australia, Australia

ABSTRACT

Decision-making is a process of choosing among alternative courses of action for solving complicated problems where multi-criteria objectives are involved. The past few years have witnessed a growing recognition of soft computing (SC) (Zadeh, 1998) technologies that underlie the conception, design, and utilization of intelligent systems. In this chapter, we present different SC paradigms involving an artificial neural network (Zurada, 1992) trained by using the scaled conjugate gradient algorithm (Moller, 1993), two different fuzzy inference methods (Abraham, 2001) optimised by using neural network learning/evolutionary algorithms (Fogel, 1999), and regression trees (Breiman, Friedman, Olshen, & Stone, 1984) for developing intelligent decision support systems (Tran, Abraham, & Jain, 2004). We demonstrate the efficiency of the different algorithms by developing a decision support system for a tactical air combat environment (TACE) (Tran & Zahid, 2000). Some empirical comparisons between the different algorithms are also provided.

INTRODUCTION

Several decision support systems have been developed in various fields including medical diagnosis (Adibi, Ghoreishi, Fahimi, & Maleki, 1993), business management, control system (Takagi & Sugeno, 1983), command and control of defence and air traffic control (Chappel, 1992), and so on. Usually previous experience or expert knowledge is often used to design decision support systems. The task becomes interesting when no prior knowledge is available. The need for an intelligent mechanism for decision support comes from the well-known limits of human knowledge processing. It has been noticed that the need for support for human decision-makers is due to four kinds of limits: cognitive, economic, time, and competitive demands (Holsapple & Whinston, 1996). Several artificial intelligence techniques have been explored to construct adaptive decision support systems. A framework that could capture imprecision, uncertainty, learn from the data/information, and continuously optimise the solution by providing interpretable decision rules, would be the ideal technique. Several adaptive learning frameworks for constructing intelligent decision support systems have been proposed (Cattral, Oppacher, & Deogo, 1999; Hung, 1993; Jagielska, 1998; Tran, Jain, & Abraham, 2002b). Figure 1 summarizes the basic functional aspects of a decision support system. A database is created from the available data and human knowledge. The learning process then builds up the decision rules. The developed rules are further fine-tuned, depending upon the quality of the solution, using a supervised learning process.

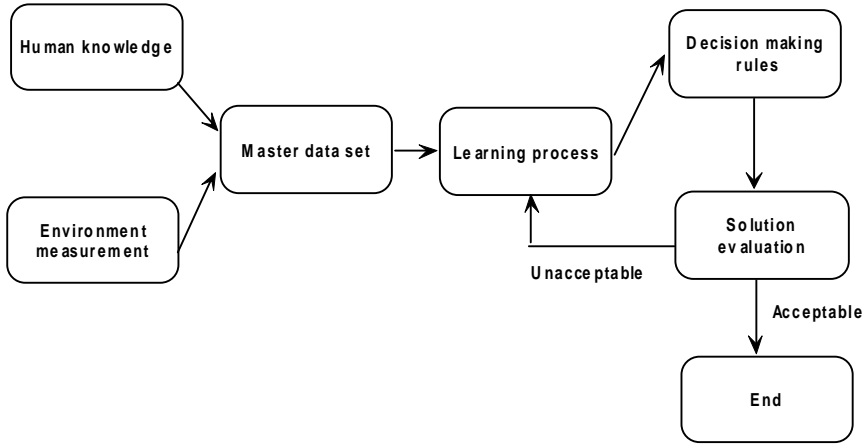
To develop an intelligent decision support system, we need a holistic view on the various tasks to be carried out including data management and knowledge management (reasoning techniques). The focus of this chapter is knowledge management (Tran & Zahid, 2000), which consists of facts and inference rules used for reasoning (Abraham, 2000).

Fuzzy logic (Zadeh, 1973), when applied to decision support systems, provides formal methodology to capture valid patterns of reasoning about uncertainty. Artificial neural networks (ANNs) are popularly known as black-box function approximators. Recent research work shows the capabilities of rule extraction from a trained network positions neuro-computing as a good decision support tool (Setiono, 2000; Setiono, Leow, & Zurada, 2002). Recently evolutionary computation (EC) (Fogel, 1999) has been successful as a powerful global optimisation tool due to the success in several problem domains (Abraham, 2002; Cortes, Larrañeta, Onieva, García, & Caraballo, 2001; Ponnuswamy, Amin, Jha, & Castañon, 1997; Tan & Li, 2001; Tan, Yu, Heng, & Lee, 2003). EC works by simulating evolution on a computer by iterative generation and alteration processes, operating on a set of candidate solutions that form a population. Due to the complementarity of neural networks, fuzzy inference systems, and evolutionary computation, the recent trend is to fuse various systems to form a more powerful integrated system, to overcome their individual weakness.

Decision trees (Breiman et al., 1984) have emerged as a powerful machine-learning technique due to a simple, apparent, and fast reasoning process. Decision trees can be related to artificial neural networks by mapping them into a class of ANNs or entropy nets with far fewer connections.

In the next section, we present the complexity of the tactical air combat decision support system (TACDSS) (Tran, Abraham, & Jain, 2002c), followed by some theoretical foundation on neural networks, fuzzy inference systems, and decision trees in the

Figure 1. Database learning framework for decision support system



following section. We then present different adaptation procedures for optimising fuzzy inference systems. A Takagi-Sugeno (Takagi & Sugeno, 1983; Sugeno, 1985) and Mamdani-Assilian (Mamdani & Assilian, 1975) fuzzy inference system learned by using neural network learning techniques and evolutionary computation is discussed. Experimental results using the different connectionist paradigms follow. Detailed discussions of these results are presented in the last section, and conclusions are drawn.

TACTICAL AIR COMBAT DECISION SUPPORT SYSTEM

Implementation of a reliable decision support system involves two important factors: collection and analysis of prior information, and the evaluation of the solution. The data could be an image or a pattern, real number, binary code, or natural language text data, depending on the objects of the problem environment. An object of the decision problem is also known as the decision factor. These objects can be expressed mathematically in the decision problem domain as a universal set, where the decision factor is a set and the decision data is an element of this set. The decision factor is a sub-set of the decision problem. If we call the decision problem (DP) as X and the decision factor (DF) as “ A ”, then the decision data (DD) could be labelled as “ a ”. Suppose the set A has members a_1, a_2, \dots, a_n then it can be denoted by $A = \{a_1, a_2, \dots, a_n\}$ or can be written as:

$$A = \{a_i \mid i \in R_n\} \quad (1)$$

where i is called the *set index*, the symbol “ \mid ” is read as “such that” and R_n is the set of n real numbers. A sub-set “ A ” of X , denoted $A \subseteq X$, is a set of elements that is contained within the universal set X . For optimal decision-making, the system should be able to

adaptively process the information provided by words or any natural language description of the problem environment.

To illustrate the proposed approach, we consider a case study based on a tactical environment problem. We aim to develop an environment decision support system for a pilot or mission commander in tactical air combat. We will attempt to present the complexity of the problem with some typical scenarios. In Figure 2, the Airborne Early Warning and Control (AEW&C) is performing surveillance in a particular area of operation. It has two Hornets (F/A-18s) under its control at the ground base shown as “+” in the left corner of Figure 2. An air-to-air fuel tanker (KB707) “□” is on station — the location and status of which are known to the AEW&C. One of the Hornets is on patrol in the area of Combat Air Patrol (CAP). Sometime later, the AEW&C on-board sensors detect hostile aircraft(s) shown as “O”. When the hostile aircraft enter the surveillance region (shown as a dashed circle), the mission system software is able to identify the enemy aircraft and estimate their distance from the Hornets in the ground base or in the CAP.

The mission operator has few options to make a decision on the allocation of Hornets to intercept the enemy aircraft:

- Send the Hornet directly to the spotted area and intercept,
- Call the Hornet in the area back to ground base or send another Hornet from the ground base.
- Call the Hornet in the area for refuel before intercepting the enemy aircraft.

The mission operator will base his/her decisions on a number of factors, such as:

- Fuel reserve and weapon status of the Hornet in the area,
- Interrupt time of Hornets in the ground base or at the CAP to stop the hostile,
- The speed of the enemy fighter aircraft and the type of weapons it possesses.

Figure 2. A typical air combat scenario




Table 1. Decision factors for the tactical air combat

Fuel reserve	Intercept time	Weapon status	Danger situation	Evaluation plan
Full	Fast	Sufficient	Very dangerous	Good
Half	Normal	Enough	Dangerous	Acceptable
Low	Slow	Insufficient	Endangered	Bad

From the above scenario, it is evident that there are important decision factors of the tactical environment that might directly affect the air combat decision. For demonstrating our proposed approach, we will simplify the problem by handling only a few important decision factors such as “fuel status”, “interrupt time” (Hornets in the ground base and in the area of CAP), “weapon possession status”, and “situation awareness” (Table 1). The developed tactical air combat decision rules (Abraham & Jain, 2002c) should be able to incorporate all the above-mentioned decision factors.

Knowledge of Tactical Air Combat Environment

How can human knowledge be extracted to a database? Very often people express knowledge as natural (spoken) language or using letters or symbolic terms. The human knowledge can be analysed and converted into an information table. There are several methods to extract human knowledge. Some researchers use cognitive work analysis (CWA) (Sanderson, 1998); others use cognitive task analysis (CTA) (Mitallo, 1998). CWA is a technique used to analyse, design, and evaluate human computer interactive systems. CTA is a method used to identify cognitive skills and mental demands, and needs to perform these tasks proficiently. CTA focuses on describing the representation of the cognitive elements that define goal generation and decision making. It is a reliable method to extract human knowledge because it is based on observations or an interview. We have used the CTA technique to set up the expert knowledge base for building the complete decision support system. For the TACE discussed previously, we have four decision factors that could affect the final decision options of “Hornet in the CAP” or “Hornet at the ground base”. These are: “fuel status” being the quantity of fuel available to perform the intercept, the “weapon possession status” presenting the state of available weapons inside the Hornet, the “interrupt time” which is required for the Hornet to fly and interrupt the hostile, and the “danger situation” providing information whether the aircraft is friendly or hostile.

Each of the above-mentioned factors has a different range of units, these being the fuel (0 to 1000 litres), interrupt time (0 to 60 minutes), weapon status (0 to 100 %), and the danger situation (0 to 10 points). The following are two important decision selection rules, which were formulated using expert knowledge:

- The decision selection will have a small value if the fuel is too low, the interrupt time is too long, the Hornet has low weapon status, and the Friend-Or-Enemy/Foe danger is high.

Table 2. Some prior knowledge of the TACE

Fuel status (litres)	Interrupt time (minutes)	Weapon status (percent)	Danger situation (points)	Decision selection (points)
0	60	0	10	0
100	55	15	8	1
200	50	25	7	2
300	40	30	5	3
400	35	40	4.5	4
500	30	60	4	5
600	25	70	3	6
700	15	85	2	7
800	10	90	1.5	8
900	5	96	1	9
1000	1	100	0	10

- The decision selection will have a high value if the fuel reserve is full, the interrupt time is fast enough, the Hornet has high weapon status, and the FOE danger is low.

In TACE, decision-making is always based on all states of all the decision factors. However, sometimes a mission operator/commander can make a decision based on an important factor, such as: The fuel reserve of the Hornet is too low (due to high fuel use), the enemy has more powerful weapons, and the quality and quantity of enemy aircraft. Table 2 shows the decision score at each stage of the TACE.

SOFT COMPUTING AND DECISION TREES

Soft computing paradigms can be used to construct new generation intelligent hybrid systems consisting of artificial neural networks, fuzzy inference systems, approximate reasoning, and derivative free optimisation techniques. It is well known that the intelligent systems which provide human-like expertise such as domain knowledge, uncertain reasoning, and adaptation to a noisy and time-varying environment, are important in tackling real-world problems.

Artificial Neural Networks

Artificial neural networks have been developed as generalisations of mathematical models of biological nervous systems. A neural network is characterised by the network architecture, the connection strength between pairs of neurons (weights), node properties, and update rules. The update or learning rules control the weights and/or states of the processing elements (neurons). Normally, an objective function is defined that represents the complete status of the network, and its set of minima corresponds to different stable states (Zurada, 1992). It can learn by adapting its weights to changes in the surrounding environment, can handle imprecise information, and generalise from known tasks to unknown ones. The network is initially randomised to avoid imposing any

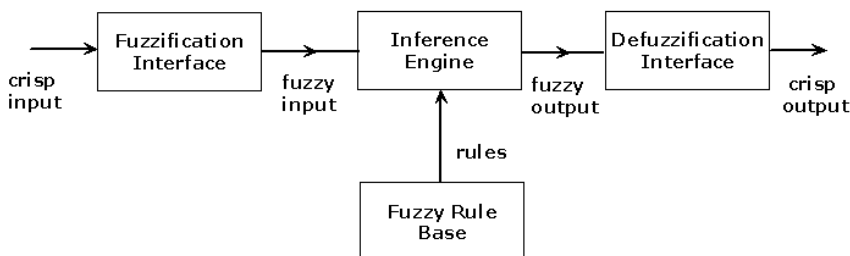
of our own prejudices about an application of interest. The training patterns can be thought of as a set of ordered pairs $\{(x_1, y_1), (x_2, y_2), \dots, (x_p, y_p)\}$ where x_i represents an input pattern and y_i represents the output pattern vector associated with the input vector x_i .

A valuable property of neural networks is that of generalisation, whereby a trained neural network is able to provide a correct matching in the form of output data for a set of previously-unseen input data. Learning typically occurs through training, where the training algorithm iteratively adjusts the connection weights (synapses). In the conjugate gradient algorithm (CGA), a search is performed along conjugate directions, which produces generally faster convergence than steepest descent directions. A search is made along the conjugate gradient direction to determine the step size, which will minimise the performance function along that line. A line search is performed to determine the optimal distance to move along the current search direction. Then the next search direction is determined so that it is conjugate to the previous search direction. The general procedure for determining the new search direction is to combine the new steepest descent direction with the previous search direction. An important feature of CGA is that the minimization performed in one step is not partially undone by the next, as is the case with gradient descent methods. An important drawback of CGA is the requirement of a line search, which is computationally expensive. The scaled conjugate gradient algorithm (SCGA) (Moller, 1993) was designed to avoid the time-consuming line search at each iteration, and incorporates the model-trust region approach used in the CGA Levenberg-Marquardt algorithm (Abraham, 2002).

Fuzzy Inference Systems (FIS)

Fuzzy inference systems (Zadeh, 1973) are a popular computing framework based on the concepts of fuzzy set theory, fuzzy if-then rules, and fuzzy reasoning. The basic structure of the fuzzy inference system consists of three conceptual components: a rule base, which contains a selection of fuzzy rules; a database, which defines the membership functions used in the fuzzy rule; and a reasoning mechanism, which performs the inference procedure upon the rules and given facts to derive a reasonable output or conclusion. Figure 3 shows the basic architecture of a FIS with crisp inputs and outputs implementing a non-linear mapping from its input space to its output (Cattral, Oppacher, & Deogo, 1992).

Figure 3. Fuzzy inference system block diagram



We now introduce two different fuzzy inference systems that have been widely employed in various applications. These fuzzy systems feature different consequents in their rules, and thus their aggregation and defuzzification procedures differ accordingly.

Most fuzzy systems employ the inference method proposed by Mamdani-Assilian in which the rule consequence is defined by fuzzy sets and has the following structure (Mamdani & Assilian, 1975):

$$\text{If } x \text{ is } A_1 \text{ and } y \text{ is } B_1 \text{ then } z_1 = C_1 \quad (2)$$

Takagi and Sugeno (1983) proposed an inference scheme in which the conclusion of a fuzzy rule is constituted by a weighted linear combination of the crisp inputs rather than a fuzzy set, and which has the following structure:

$$\text{If } x \text{ is } A_1 \text{ and } y \text{ is } B_1, \text{ then } z_1 = p_1 + q_1 y + r \quad (3)$$

A Takagi-Sugeno FIS usually needs a smaller number of rules, because its output is already a linear function of the inputs rather than a constant fuzzy set (Abraham, 2001).

Evolutionary Algorithms

Evolutionary algorithms (EAs) are population-based adaptive methods, which may be used to solve optimisation problems, based on the genetic processes of biological organisms (Fogel, 1999; Tan et al., 2003). Over many generations, natural populations evolve according to the principles of natural selection and “survival-of-the-fittest”, first clearly stated by Charles Darwin in “On the Origin of Species”. By mimicking this process, EAs are able to “evolve” solutions to real-world problems, if they have been suitably encoded. The procedure may be written as the difference equation (Fogel, 1999):

$$x[t+1] = s(v(x[t])) \quad (4)$$

Figure 4. Evolutionary algorithm pseudo code

1. Generate the initial population $P(0)$ at random and set $i=0$;
2. Repeat until the number of iterations or time has been reached or the population has converged.
 - a. Evaluate the fitness of each individual in $P(i)$
 - b. Select parents from $P(i)$ based on their fitness in $P(i)$
 - c. Apply reproduction operators to the parents and produce offspring, the next generation, $P(i+1)$ is obtained from the offspring and possibly parents.

where $x(t)$ is the population at time t , v is a random operator, and s is the selection operator. The algorithm is illustrated in Figure 4.

A conventional fuzzy controller makes use of a model of the expert who is in a position to specify the most important properties of the process. Expert knowledge is often the main source to design the fuzzy inference systems. According to the performance measure of the problem environment, the membership functions and rule bases are to be adapted. Adaptation of fuzzy inference systems using evolutionary computation techniques has been widely explored (Abraham & Nath, 2000a, 2000b). In the following section, we will discuss how fuzzy inference systems could be adapted using neural network learning techniques.

Neuro-Fuzzy Computing

Neuro-fuzzy (NF) (Abraham, 2001) computing is a popular framework for solving complex problems. If we have knowledge expressed in linguistic rules, we can build a FIS; if we have data, or can learn from a simulation (training), we can use ANNs. For building a FIS, we have to specify the fuzzy sets, fuzzy operators, and the knowledge base. Similarly, for constructing an ANN for an application, the user needs to specify the architecture and learning algorithm. An analysis reveals that the drawbacks pertaining to these approaches seem complementary and, therefore, it is natural to consider building an integrated system combining the concepts. While the learning capability is an advantage from the viewpoint of FIS, the formation of a linguistic rule base will be advantageous from the viewpoint of ANN (Abraham, 2001).

In a fused NF architecture, ANN learning algorithms are used to determine the parameters of the FIS. Fused NF systems share data structures and knowledge representations. A common way to apply a learning algorithm to a fuzzy system is to represent it in a special ANN-like architecture. However, the conventional ANN learning algorithm (gradient descent) cannot be applied directly to such a system as the functions used in the inference process are usually non-differentiable. This problem can be tackled by using differentiable functions in the inference system or by not using the standard neural learning algorithm. Two neuro-fuzzy learning paradigms are presented later in this chapter.

Classification and Regression Trees

Tree-based models are useful for both classification and regression problems. In these problems, there is a set of classification or predictor variables (X_i) and a dependent variable (Y). The X_i variables may be a mixture of nominal and/or ordinal scales (or code intervals of equal-interval scale) and Y may be a quantitative or a qualitative (in other words, nominal or categorical) variable (Breiman et al., 1984; Steinberg & Colla, 1995).

The classification and regression trees (CART) methodology is technically known as binary recursive partitioning. The process is binary because parent nodes are always split into exactly two child nodes, and recursive because the process can be repeated by treating each child node as a parent. The key elements of a CART analysis are a set of rules for splitting each node in a tree:

- deciding when a tree is complete, and
- assigning each terminal node to a class outcome (or predicted value for regression)

CART is the most advanced decision tree technology for data analysis, pre-processing, and predictive modelling. CART is a robust data-analysis tool that automatically searches for important patterns and relationships and quickly uncovers hidden structure even in highly complex data. CARTs binary decision trees are more sparing with data and detect more structure before further splitting is impossible or stopped. Splitting is impossible if only one case remains in a particular node, or if all the cases in that node are exact copies of each other (on predictor variables). CART also allows splitting to be stopped for several other reasons, including that a node has too few cases (Steinberg & Colla, 1995).

Once a terminal node is found, we must decide how to classify all cases falling within it. One simple criterion is the plurality rule: The group with the greatest representation determines the class assignment. CART goes a step further: Because each node has the potential for being a terminal node, a class assignment is made for every node whether it is terminal or not. The rules of class assignment can be modified from simple plurality to account for the costs of making a mistake in classification and to adjust for over- or under-sampling from certain classes.


A common technique among the first generation of tree classifiers was to continue splitting nodes (growing the tree) until some goodness-of-split criterion failed to be met. When the quality of a particular split fell below a certain threshold, the tree was not grown further along that branch. When all branches from the root reached terminal nodes, the tree was considered complete. Once a maximal tree is generated, it examines smaller trees obtained by pruning away branches of the maximal tree. Once the maximal tree is grown and a set of sub-trees is derived from it, CART determines the best tree by testing for error rates or costs. With sufficient data, the simplest method is to divide the sample into learning and test sub-samples. The learning sample is used to grow an overly large tree. The test sample is then used to estimate the rate at which cases are misclassified (possibly adjusted by misclassification costs). The misclassification error rate is calculated for the largest tree and also for every sub-tree.

The best sub-tree is the one with the lowest or near-lowest cost, which may be a relatively small tree. Cross validation is used if data are insufficient for a separate test sample. In the search for patterns in databases, it is essential to avoid the trap of over-fitting or finding patterns that apply only to the training data. CARTs embedded test disciplines ensure that the patterns found will hold up when applied to new data. Further, the testing and selection of the optimal tree are an integral part of the CART algorithm. CART handles missing values in the database by substituting surrogate splitters, which are back-up rules that closely mimic the action of primary splitting rules. The surrogate splitter contains information that is typically similar to what would be found in the primary splitter (Steinberg & Colla, 1995).

TACDSS ADAPTATION USING TAKAGI-SUGENO FIS

We used the adaptive network-based fuzzy inference system (ANFIS) framework (Jang, 1992) to develop the TACDSS based on a Takagi-Sugeno fuzzy inference system. The six-layered architecture of ANFIS is depicted in Figure 5.

Figure 5. ANFIS architecture



Suppose there are two input linguistic variables (ILV) X and Y and each ILV has three membership functions (MF) A_1, A_2 and A_3 and B_1, B_2 and B_3 respectively, then a Takagi-Sugeno-type fuzzy *if-then* rule could be set up as:

$$\text{Rule}_i : \text{If } X \text{ is } A_i \text{ and } Y \text{ is } B_i \text{ then } f_i = p_i X + q_i Y + r_i \quad (5)$$

where i is an index $i = 1, 2, \dots, n$ and p, q and r are the linear parameters.

Some layers of ANFIS have the same number of nodes, and nodes in the same layer have similar functions. Output of nodes in layer- l is denoted as $O_{l,p}$ where l is the layer number and i is neuron number of the next layer. The function of each layer is described as follows:

- **Layer 1**

The outputs of this layer is the input values of the ANFIS

$$O_{1,x} = x$$

$$O_{1,y} = y \quad (6)$$

For TACDSS the four inputs are “fuel status”, “weapons inventory levels”, “time intercept”, and the “danger situation”.

- **Layer 2**

The output of nodes in this layer is presented as $O_{l,ip,i}$, where ip is the ILV and m is the degree of membership function of a particular MF.

$$O_{2,x,i} = m_{Ai(x)} \text{ or } O_{2,y,i} = m_{Bi(y)} \text{ for } i = 1, 2, \text{ and } 3 \quad (7)$$

With three MFs for each input variable, “fuel status” has three membership functions: *full*, *half*, and *low*, “time intercept” has *fast*, *normal*, and *slow*, “weapon status” has *sufficient*, *enough*, and *insufficient*, and the “danger situation” has *very dangerous*, *dangerous*, and *endangered*.

- **Layer 3**

The output of nodes in this layer is the product of all the incoming signals, denoted by:

$$O_{3,n} = W_n = m_{Ai}(x) \times m_{Bi}(y) \quad (8)$$

where $i = 1, 2, \text{ and } 3$, and n is the number of the fuzzy rule. In general, any T-norm operator will perform the fuzzy ‘AND’ operation in this layer. With four ILV and three MFs for each input variable, the TACDSS will have 81 ($3^4 = 81$) fuzzy *if-then* rules.

- **Layer 4**

The nodes in this layer calculate the ratio of the i^{th} fuzzy rule firing strength (RFS) to the sum of all RFS.

$$O_{4,n} = \overline{w_n} = \frac{w_n}{\sum_{n=1}^{81} w_n} \text{ where } n = 1, 2, \dots, 81 \quad (9)$$

The number of nodes in this layer is the same as the number of nodes in layer-3. The outputs of this layer are also called normalized firing strengths.

- **Layer 5**

The nodes in this layer are adaptive, defined as:

$$O_{5,n} = \overline{w_n} f_n = \overline{w_n} (p_n x + q_n y + r_n) \quad (10)$$

where p_n , q_n , r_n are the rule *consequent parameters*. This layer also has the same number of nodes as layer-4 (81 numbers).

- **Layer 6**

The single node in this layer is responsible for the defuzzification process, using the centre-of-gravity technique to compute the overall output as the summation of all the incoming signals:

$$O_{6,1} = \sum_{n=1}^{81} \overline{w_n} f_n = \frac{\sum_{n=1}^{81} w_n f_n}{\sum_{n=1}^{81} w_n} \quad (11)$$

ANFIS makes use of a mixture of back-propagation to learn the premise parameters and least mean square estimation to determine the consequent parameters. Each step in the learning procedure comprises two parts: In the first part, the input patterns are propagated, and the optimal conclusion parameters are estimated by an iterative least mean square procedure, while the antecedent parameters (membership functions) are assumed to be fixed for the current cycle through the training set. In the second part, the patterns are propagated again, and in this epoch, back-propagation is used to modify the antecedent parameters, while the conclusion parameters remain fixed. This procedure is then iterated, as follows (Jang, 1992):

$$\begin{aligned}
 \text{ANFIS output } f = O_{6,1} &= \frac{w_1}{\sum_n w_n} f_1 + \frac{w_2}{\sum_n w_n} f_2 + \dots + \frac{w_n}{\sum_n w_n} f_n \\
 &= \bar{w}_1 (p_1 x + q_1 y + r_1) + \bar{w}_2 (p_2 x + q_2 y + r_2) + \dots + (p_n x + q_n y + r_n) \\
 &= (x)p_1 + (y)q_1 + r_1 + (x)p_2 + (y)q_2 + r_2 + \dots + (x)p_n + (y)q_n + r_n \quad (12)
 \end{aligned}$$

where n is the number of nodes in layer 5. From this, the output can be rewritten as

$$f = F(i, S) \quad (13)$$

where F is a function, i is the vector of input variables, and S is a set of total parameters of consequent of the n^{th} fuzzy rule. If there exists a composite function H such that $H \oplus F$ is linear in some elements of S , then these elements can be identified by the least square method. If the parameter set is divided into two sets S_1 and S_2 , defined as:

$$S = S_1 \oplus S_2 \quad (14)$$

where \oplus represents direct sum and \circ is the product rule, such that $H \circ F$ is linear in the elements of S_2 , the function f can be represented as:

$$H(f) = H \oplus F(I, S) \quad (15)$$

Given values of S_1 , the S training data can be substituted into equation 15. $H(f)$ can be written as the matrix equation of $AX = Y$, where X is an unknown vector whose elements are parameters in S_2 .

If $|S_2| = M$ (M being the number of linear parameters), then the dimensions of matrices A , X and Y are PM , MI and PI , respectively. This is a standard linear least-squares problem and the best solution of X that minimizes $\|AX - Y\|^2$ is the least square estimate (LSE) X^*

$$X^* = (A^T A)^{-1} A^T Y \quad (16)$$

where A^T is the transpose of A , $(A^T A)^{-1} A^T$ is the pseudo inverse of A if $A^T A$ is a non-singular. Let the i^{th} row vector of matrix A be a and the i^{th} element of Y be y , then X can be calculated as:

$$X_{i+1} = X_i + S_{i+1} a_{i+1} (y - y - aX_i) \quad (17)$$

$$S_{i+1} = S_i - \frac{S_i a_{i+1} y_{i+1}^T - S_i}{1 + a_{i+1}^T S_i a_{i+1}}, \quad I = 0, 1, \dots, P-1 \quad (18)$$

The LSE X^* is equal to X_p . The initial conditions of X_{i+1} and S_{i+1} are $X_0 = 0$ and $S_0 = gI$, where g is a positive large number and I is the identity matrix of dimension $M \times M$.

When hybrid learning is applied in batch mode, each epoch is composed of a forward pass and a backward pass. In the forward pass, the node output I of each layer is calculated until the corresponding matrices A and Y are obtained. The parameters of S_2 are identified by the pseudo inverse equation as mentioned above. After the parameters of S_2 are obtained, the process will compute the error measure for each training data pair. In the backward pass, the error signals (the derivatives of the error measure with respect to each node output) propagates from the output to the input end. At the end of the backward pass, the parameter S_i is updated by the steepest descent method as follows:

$$a = -\eta \frac{\partial E}{\partial \alpha} \quad (19)$$

where a is a generic parameter and η is the learning rate and E is an error measure.

$$\eta = \frac{k}{\sqrt{\sum \alpha \left(\frac{\partial E}{\partial \alpha} \right)^2}} \quad (20)$$

where k is the step size.

For the given fixed values of parameters in S_1 , the parameters in S_2 are guaranteed to be global optimum points in the S_2 parameters space due to the choice of the squared error measure. This hybrid learning method can decrease the dimension of the search space using the steepest descent method, and can reduce the time needed to reach convergence. The step size k will influence the speed of convergence. Observation shows that if k is small, the gradient method will closely approximate the gradient path; convergence will be slow since the gradient is being calculated many times. If the step size k is large, convergence will initially be very fast. Based on these observations, the step size k is updated by the following two heuristics (rules) (Jang, 1992):

If E undergoes four continuous reductions, then increase k by 10%, and

If E undergoes continuous combinations of increase and decrease, then reduce k by 10%.

TACDSS ADAPTATION USING MAMDANI FIS

We have made use of the fuzzy neural network (FuNN) framework (Kasabov, Kim & Gray, 1996) for learning the Mamdani-Assilian fuzzy inference method. A functional block diagram of the FuNN model is depicted in Figure 6 (Kasabov, 1996); it consists of two phases of learning.

The first phase is the structure learning (*if-then* rules) using the knowledge acquisition module. The second phase is the parameter learning for tuning membership functions to achieve a desired level of performance. FuNN uses a gradient descent learning algorithm to fine-tune the parameters of the fuzzy membership functions. In the connectionist structure, the input and output nodes represent the input states and output control-decision signals, respectively, while in the hidden layers, there are nodes functioning as quantification of membership functions (MFs) and *if-then* rules. We used the simple and straightforward method proposed by Wang and Mendel (1992) for generating fuzzy rules from numerical input-output training data. The task here is to generate a set of fuzzy rules from the desired input-output pairs and then use these fuzzy rules to determine the complete structure of the TACDSS.

Suppose we are given the following set of desired input (x_1, x_2) and output (y) data pairs (x_1, x_2, y): (0.6, 0.2; 0.2), (0.4, 0.3; 0.4). In TACDSS, the input variable “fuel reserve” has a degree of 0.8 in *half*, a degree of 0.2 in *full*. Similarly, the input variable “time intercept” has a degree of 0.6 in *empty* and 0.3 in *normal*. Secondly, assign x_1^i, x_2^i , and y^i to a region that has maximum degree. Finally, obtain one rule from one pair of desired input-output data, for example:

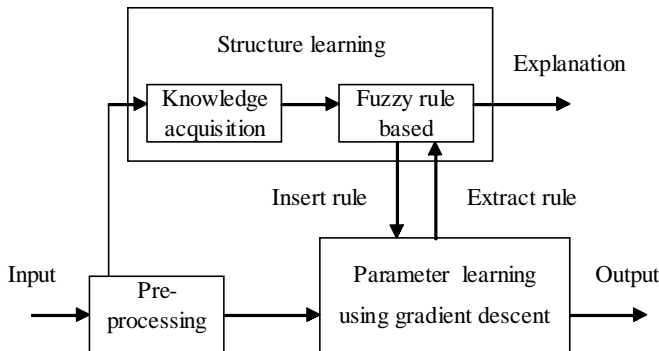
$$(x_1^1, x_2^1, y^1) \Rightarrow [x_1^1 (0.8 \text{ in } half), x_2^1 (0.2 \text{ in } fast), y^1 (0.6 \text{ in } acceptable)],$$

$$R_1: \text{ if } x_1 \text{ is } half \text{ and } x_2 \text{ is } fast, \text{ then } y \text{ is } acceptable \quad (21)$$

$$(x_1^2, x_2^2, y^2) \Rightarrow [x_1 (0.8 \text{ in } half), x_2 (0.6 \text{ in } normal), y^2 (0.8 \text{ in } acceptable)],$$

$$R_2: \text{ if } x_1 \text{ is } half \text{ and } x_2 \text{ is } normal, \text{ then } y \text{ is } acceptable \quad (22)$$

Figure 6. A general schematic of the hybrid fuzzy neural network



Assign a degree to each rule. To resolve a possible conflict problem, that is, rules having the same antecedent but a different consequent, and to reduce the number of rules, we assign a degree to each rule generated from data pairs and accept only the rule from a conflict group that has a maximum degree. In other words, this step is performed to delete redundant rules, and therefore obtain a concise fuzzy rule base. The following product strategy is used to assign a degree to each rule. The degree of the rule is denoted by:

$$R_i : \text{if } x_1 \text{ is } A \text{ and } x_2 \text{ is } B, \text{ then } y \text{ is } C(w_i) \quad (23)$$

The rule weight is defined as:

$$w_i = m_A(x_1)m_B(x_2)m_C(y) \quad (24)$$

For example in the TACE, R_1 has a degree of

$$W_1 = m_{half}(x_1) m_{fast}(x_2) m_{acceptable}(y) = 0.8 \times 0.2 \times 0.6 = 0.096 \quad (25)$$

and R_2 has a degree of

$$W_2 = m_{half}(x_1) m_{normal}(x_2) m_{acceptable}(y) = 0.8 \times 0.6 \times 0.8 = 0.384 \quad (26)$$

Note that if two or more generated fuzzy rules have the same preconditions and consequents, then the rule that has maximum degree is used. In this way, assigning the degree to each rule, the fuzzy rule base can be adapted or updated by the relative weighting strategy: The more task-related the rule becomes, the more weight degree the rule gains. As a result, not only is the conflict problem resolved, but also the number of rules is reduced significantly. After the structure-learning phase (*if-then* rules), the whole network structure is established, and the network enters the second learning phase to optimally adjust the parameters of the membership functions using a gradient descent learning algorithm to minimise the error function:


$$E = \frac{1}{2} \sum_x \sum_{l=1}^q (d_l - y_l)^2 \quad (27)$$

where d and y are the target and actual outputs for an input x . This approach is very similar to the *MF* parameter tuning in ANFIS.

Membership Function Parameter Optimisation Using EAs

We have investigated the usage of evolutionary algorithms (EAs) to optimise the number of rules and fine-tune the membership functions (Tran, Jain, & Abraham, 2002a). Given that the optimisation of fuzzy membership functions may involve many changes to many different functions, and that a change to one function may affect others, the large possible solution space for this problem is a natural candidate for an EA-based approach. This has already been investigated in Mang, Lan, and Zhang (1995), and has been shown

Figure 7. The chromosome of the centres of input and output MF's



to be more effective than manual alteration. A similar approach has been taken to optimise membership function parameters. A simple way is to represent only the parameter showing the centre of MFs to speed up the adaptation process and to reduce spurious local minima over the centre and width.

The EA module for adapting FuNN is designed as a stand-alone system for optimising the MFs if the rules are already available. Both antecedent and consequent MFs are optimised. Chromosomes are represented as strings of floating-point numbers rather than strings of bits. In addition, mutation of a gene is implemented as a re-initialisation, rather than an alteration of the existing allegation. Figure 7 shows the chromosome structure, including the input and output MF parameters. One point crossover is used for the chromosome reproduction.

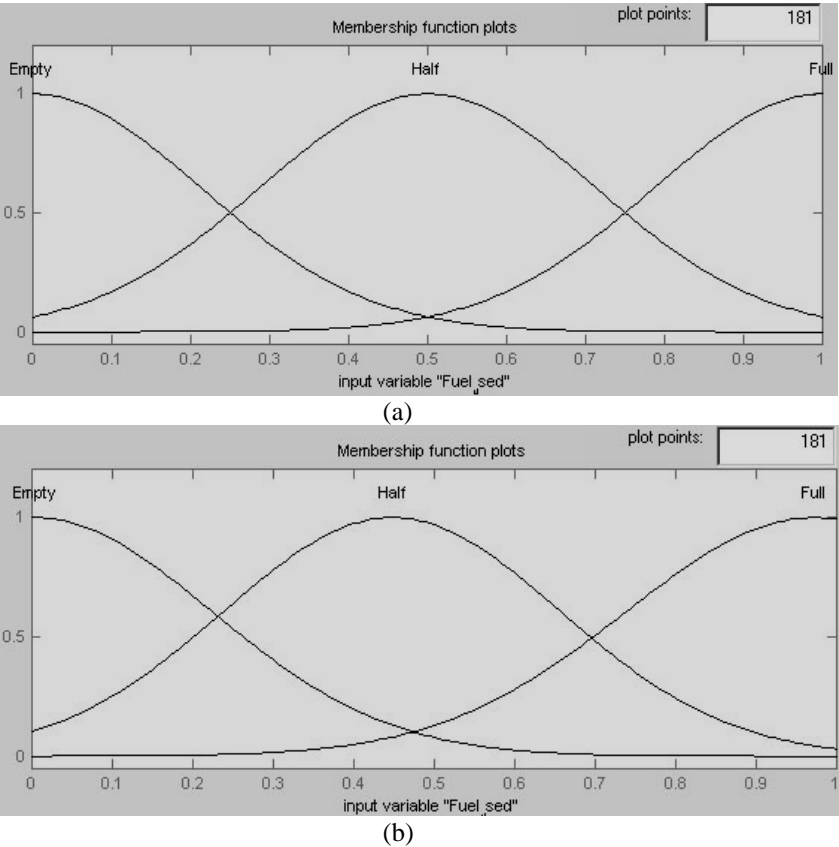
EXPERIMENTAL RESULTS FOR DEVELOPING THE TACDSS

Our master data set comprised 1000 numbers. To avoid any bias on the data, we randomly created two training sets (Dataset A - 90% and Dataset B - 80%) and test data (10% and 20 %) from the master dataset. All experiments were repeated three times and the average errors are reported here.

Takagi-Sugeno Fuzzy Inference System

In addition to the development of the Takagi-Sugeno FIS, we also investigated the behaviour of TACDSS for different membership functions (shape and quantity per ILV). We also explored the importance of different learning methods for fine-tuning the rule antecedents and consequents. Keeping the consequent parameters constant, we fine-tuned the membership functions alone using the gradient descent technique (back-propagation). Further, we used the hybrid learning method wherein the consequent parameters were also adjusted according to the least squares algorithm. Even though back-propagation is faster than the hybrid technique, learning error and decision scores were better for the latter. We used three Gaussian MFs for each ILV. Figure 8 shows the three MFs for the “fuel reserve” ILV before and after training. The fuzzy rule consequent parameters before training was set to zero, and the parameters were learned using the hybrid learning approach.

Figure 8. Membership function of the “fuel reserve” ILV (a) before and (b) after learning



Comparison of the Shape of Membership Functions of FIS

In this section, we demonstrate the importance of the shape of membership functions. We used the hybrid-learning technique and each ILV had three MFs. Table 3 shows the convergence of the training RMSE during the 15 epoch learning using four different membership functions for 90% and 80% training data. Eighty-one fuzzy *if-then* rules were created initially using a grid-partitioning algorithm. We considered Generalised bell, Gaussian, trapezoidal, and isosceles triangular membership functions. Figure 9 illustrates the training convergence curve for different MFs.

As is evident from Table 3 and Figure 9, the lowest training and test error was obtained using a Gaussian MF.

Figure 9. Effect on training error for the different membership functions




Table 3. Learning performance showing the effect of the shape of MF

Root Mean Squared Error (E- 05)								
Epochs	Gaussian		G-bell		Trapezoidal		Triangular	
	Data A	Data B	Data A	Data B	Data A	Data B	Data A	Data B
1	1.406	1.305	1.706	1.581	2.459	2.314	0.9370	0.8610
2	1.372	1.274	1.652	1.537	2.457	2.285	1.789	1.695
3	1.347	1.249	1.612	1.505	2.546	2.441	1.789	1.695
4	1.328	1.230	1.586	1.483	2.546	2.441	1.789	1.695
5	1.312	1.214	1.571	1.471	2.546	2.441	1.789	1.695
6	1.300	1.199	1.565	1.466	2.546	2.441	1.789	1.695
7	1.288	1.186	1.564	1.465	2.546	2.441	1.789	1.695
8	1.277	1.173	1.565	1.464	2.546	2.441	1.789	1.695
9	1.265	1.160	1.565	1.459	2.546	2.441	1.789	1.695
10	1.254	1.148	1.565	1.448	2.546	2.441	1.789	1.695
11	1.243	1.138	1.565	1.431	2.546	2.441	1.789	1.695
12	1.236	1.132	1.565	1.409	2.546	2.441	1.789	1.695
13	1.234	1.132	1.565	1.384	2.546	2.441	1.789	1.695
14	1.238	1.138	1.565	1.355	2.546	2.441	1.789	1.695
Test RMSE	1.44	1.22	1.78	1.36	2.661	2.910	1.8583	1.8584

Mamdani Fuzzy Inference System

We used FuzzyCOPE (Watts, Woodford, & Kasabov, 1999) to investigate the tuning of membership functions using back-propagation and evolutionary algorithms. The learning rate and momentum were set at 0.5 and 0.3 respectively, for 10 epochs. We obtained training RMSEs of 0.2865 (Data A) and 0.2894 (Data B). We further improved the training performance using evolutionary algorithms. The following settings were used for the evolutionary algorithm parameters:

Population size = 50
Number of generations = 100
Mutation rate = 0.01

We used the tournament selection strategy, and Figure 10 illustrates the learning convergence during the 100 generations for Datasets A and B. Fifty-four fuzzy *if-then* rules were extracted after the learning process. Table 4 summarizes the training and test performance.

Figure 10. Training convergence using evolutionary algorithms





Table 4. Training and test performance of Mamdani FIS using EA's

Root Mean Squared Error (RMSE)			
Data A		Data B	
Training	Test	Training	Test
0.0548	0.0746	0.0567	0.0612

Figure 11. Neural network training using SCGA*Table 5. Training and test performance of neural networks versus decision trees*

	Data A		Data B	
	Training	Testing	Training	Testing
	RMSE			
CART	0.00239	0.00319	0.00227	0.00314
Neural Network	0.00105	0.00095	0.00041	0.00062

Artificial Neural Networks

We used 30 hidden neurons for Data A and 32 hidden neurons for Data B. We used a trial-and-error approach to finalize the architecture of the neural network. We used the scaled conjugate gradient algorithm to develop the TACDSS. Training was terminated after 1000 epochs. Figure 11 depicts the convergence of training during 1000 epochs learning. Table 5 summarizes the training and test performance.

Classification and Adaptive Regression Trees

We used a CART simulation environment to develop the decision trees (www.salford-systems.com/products-cart.html). We selected the minimum cost tree regardless of tree size. Figures 12 and 13 illustrate the variation of error with reference to the number of terminal nodes for Datasets A and B. For Data A, the developed tree has 122 terminal

Figure 12. Dataset A - Variation of relative error versus the number of terminal nodes




Figure 13. Dataset B - Variation of relative error versus the number of terminal nodes




Figure 14. Dataset A - Developed decision tree with 122 nodes




Figure 15. Dataset B - Developed decision tree with 128 nodes





Figure 16. Test results illustrating the efficiency of the different intelligent paradigms used in developing the TACDSS



nodes as shown in Figure 14, while for Data B, the rest of the tree had 128 terminal nodes as depicted in Figure 15. Training and test performance are summarized in Table 5.

Figure 16 compares the performance of the different intelligent paradigms used in developing the TACDSS (for clarity, we have chosen only 20% of the test results for Dataset B).

DISCUSSION

The focus of this research is to create accurate and highly interpretable (using rules or tree structures) decision support systems for a tactical air combat environment problem.

Experimental results using two different datasets revealed the importance of fuzzy inference engines to construct accurate decision support systems. As expected, by providing more training data (90% of the randomly-chosen master data set), the models were able to learn and generalise more accurately. The Takagi-Sugeno fuzzy inference system has the lowest RMSE on both test datasets. Since learning involves a complicated

procedure, the training process of the Takagi-Sugeno fuzzy inference system took longer compared to the Mamdani-Assilian fuzzy inference method; hence, there is a compromise between performance and computational complexity (training time). Our experiments using different membership function shapes also reveal that the Gaussian membership function is the “optimum” shape for constructing accurate decision support systems.

Neural networks can no longer be considered as ‘black boxes’. Recent research (Setiono, 2000; Setiono, Leow, & Zurada, 2002) has revealed that it is possible to extract rules from trained neural networks. In our experiments, we used a neural network trained using the scaled conjugate gradient algorithm. Results depicted in Figure 5 also reveal with the trained neural network could not learn and generalise accurately compared with the Takagi-Sugeno fuzzy inference system. The proposed neural network outperformed both the Mamdani-Assilian fuzzy inference system and CART.

Two important features of the developed classification and regression tree are its easy interpretability and low complexity. Due to its one-pass training approach, the CART algorithm also has the lowest computational load. For Dataset A, the best results were achieved using 122 terminal nodes (relative error = 0.00014). As shown in Figure 12, when the number of terminal nodes was reduced to 14, the relative error increased to 0.016. For Dataset B, the best results could be achieved using 128 terminal nodes (relative error = 0.00010). As shown in Figure 13, when the terminal nodes were reduced to 14, the relative error increased to 0.011.

CONCLUSION

In this chapter, we have presented different soft computing and machine learning paradigms for developing a tactical air combat decision support system. The techniques explored were a Takagi-Sugeno fuzzy inference system trained by using neural network learning techniques, a Mamdani-Assilian fuzzy inference system trained by using evolutionary algorithms and neural network learning, a feed-forward neural network trained by using the scaled conjugate gradient algorithm, and classification and adaptive regression trees.

The empirical results clearly demonstrate that all these techniques are reliable and could be used for constructing more complicated decision support systems. Experiments on the two independent data sets also reveal that the techniques are not biased on the data itself. Compared to neural networks and regression trees, the Takagi-Sugeno fuzzy inference system has the lowest RMSE, and the Mamdani-Assilian fuzzy inference system has the highest RMSE. In terms of computational complexity, perhaps regression trees are best since they use a one-pass learning approach when compared to the many learning iterations required by all other considered techniques. An important advantage of the considered models is fast learning, easy interpretability (*if-then* rules for fuzzy inference systems, *m-of-n* rules from a trained neural network (Setiono, 2000) and decision trees), efficient storage and retrieval capacities, and so on. It may also be concluded that fusing different intelligent systems, knowing their strengths and weakness could help to mitigate the limitations and take advantage of the opportunities to produce more efficient decision support systems than those built with stand-alone systems.

Our future work will be directed towards optimisation of the different intelligent paradigms (Abraham, 2002), which we have already used, and also to develop new adaptive reinforcement learning systems that can update the knowledge from data, especially when no expert knowledge is available.

ACKNOWLEDGMENTS

The authors would like to thank Professor John Fulcher for the editorial comments which helped to improve the clarity of this chapter.

REFERENCES

- Abraham, A. (2001). Neuro-fuzzy systems: State-of-the-art modeling techniques. In J. Mira & A. Prieto (Eds.), *Connectionist models of neurons, learning processes, and artificial intelligence* (pp. 269-276). Berlin, Germany: Springer-Verlag.
- Abraham, A. (2002). Optimization of evolutionary neural networks using hybrid learning algorithms. *Proceedings of the IEEE International Joint Conference on Neural Networks (IJCNN'02): Vol. 3*, Honolulu, Hawaii (pp. 2797-2802). Piscataway, NJ: IEEE Press.
- Abraham, A., & Nath, B. (2000a). Evolutionary design of neuro-fuzzy systems: A generic framework. In A. Namatame, et al. (Eds.), *Proceedings of the 4th Japan-Australia Joint Workshop on Intelligent and Evolutionary Systems (JA2000 - Japan)* (pp. 106-113). National Defence Academy (Japan)/University of New South Wales (Australia).
- Abraham, A., & Nath, B. (2000b, December). Evolutionary design of fuzzy control systems: A hybrid approach. In J. L. Wang (Ed.), *Proceedings of the 6th International Conference on Control, Automation, Robotics, and Vision, (ICARCV 2000)*, Singapore.
- Abraham, A., & Nath, B. (2001). A neuro-fuzzy approach for modelling electricity demand in Victoria. *Applied Soft Computing*, 1(2), 127-138.
- Adibi, J., Ghoreishi, A., Fahimi, M., & Maleki, Z. (1993, April). Fuzzy logic information theory hybrid model for medical diagnostic expert system. *Proceedings of the 12th Southern Biomedical Engineering Conference*, Tulane University, New Orleans, LA (pp. 211-213).
- Breiman, L., Friedman, J., Olshen, R., & Stone, C. J. (1984). *Classification and regression trees*. New York: Chapman and Hall.
- Cattral R., Oppacher F., & Deogo, D. (1999, July 6-9). Rule acquisition with a genetic algorithm. *Proceedings of the Congress on Evolution Computation: Vol. 1*, Washington, DC (pp. 125-129). Piscataway, NJ: IEEE Press.
- Chappel, A. R. (1992, October 5-8). Knowledge-based reasoning in the Paladin tactical decision generation system. *Proceedings of the 11th AIAA Digital Avionics Systems Conference*, Seattle, WA (pp. 155-160).
- Cortés, P., Larrañeta, J., Onieva, L., García, J. M., & Caraballo, M. S. (2001). Genetic algorithm for planning cable telecommunication networks. *Applied Soft Computing*, 1(1), 21-33.

- Fogel, D. (1999). *Evolutionary computation: Towards a new philosophy of machine intelligence* (2nd ed.). Piscataway, NJ: IEEE Press.
- Gorzalczany, M. B. (1996, June 17-20). An idea of the application of fuzzy neural networks to medical decision support systems. *Proceedings of the IEEE International Symposium on Industrial Electronics (ISIE '96): Vol. 1*, Warsaw, Poland (pp. 398-403).
- Holland, J. H., Kaufmann, M., & Altos, L. (1986). Escaping brittleness: The possibility of general-purpose learning algorithms applied to rule-based systems. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (pp. 593-623). San Mateo, CA: Morgan Kaufmann.
- Holsapple, C. W., & Winston, A. B. (1996). *Decision support systems: A knowledge-based approach*. Minneapolis, MN: West Publishing Company.
- Hung, C. C. (1993, November). Building a neuro-fuzzy learning control system. *AI Expert*, 8(10), 40-49.
- Ichimura, T., Takano, T., & Tazaki, E. (1995, October 8-11). Reasoning and learning method for fuzzy rules using neural networks with adaptive structured genetic algorithm. *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics — Intelligent Systems for the 21st Century: Vol. 4*, Vancouver, Canada (pp. 3269-3274).
- Jagielska, I. (1998, April 21-23). Linguistic rule extraction from neural networks for descriptive data mining. *Proceedings of the 2nd Conference on Knowledge-Based Intelligent Electronic Systems — KES'98: Vol. 2*, Adelaide, South Australia (pp. 89-92). Piscataway, NJ: IEEE Press.
- Jang, R. (1992, July). *Neuro-fuzzy modeling: Architectures, analyses, and applications*. PhD Thesis, University of California, Berkeley.
- Kasabov, N. (1996). Learning fuzzy rules and approximate reasoning in fuzzy neural networks and hybrid systems. *Fuzzy Sets and Systems*, 82, 135-149.
- Kasabov, N. (2001, December). Evolving fuzzy neural networks for supervised/unsupervised on-line knowledge-based learning. *IEEE Transaction of Systems, Man, and Cybernetics, Part B — Cybernetic*, 31(6), 902-918.
- Kasabov, N., Kim, J. S., & Gray, A. R. (1996). FUNN — A fuzzy neural network architecture for adaptive learning and knowledge acquisition. *Information Sciences*, 101(3), 155-175.
- Kearney, D. A., & Tran, C. M. (1995, October 23-25). Optimal fuzzy controller design for minimum rate of change of acceleration in a steel uncoiler. *Control95 — Meeting the Challenge of Asia Pacific Growth: Vol. 2*, University of Melbourne, Australia (pp. 393-397).
- Lee, C. C. (1990). Fuzzy logic control systems: Fuzzy logic controller — Part I & II. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(2), 404-435.
- Lin, T. Y., & Cercone, N. (1997). *Rough sets and data mining: Analysis of imprecise data*. New York: Kluwer Academic.
- Mamdani, E. H., & Assilian, S. (1975). An experiment in linguistic synthesis with a fuzzy logic controller. *International Journal of Man-Machine Studies*, 7(1), 1-13.
- Mang, G., Lan, H., & Zhang, L. (1995, October 30-November 3). A genetic-base method of generating fuzzy rules and membership function by learning from examples. *Proceedings of the International Conference on Neural Information (ICONIP '95): Vol. 1*, Beijing, China (pp. 335-338).

- Militallo, L. G., & Hutton, R. J. B. (1998). Applied cognitive task analysis (ACTA): A practitioner's toolkit for understanding cognitive. *Ergonomics*, 41(11), 1618-1642.
- Moller, A. F. (1993). A scaled conjugate gradient algorithm for fast supervised learning. *Neural Networks*, 6, 525-533.
- Perneel, C., & Acheroy, M. (1994, December 12-13). Fuzzy reasoning and genetic algorithm for decision making problems in uncertain environment. *Proceedings of the Industrial Fuzzy Control and Intelligent Systems Conference/NASA joint Technology Workshop on Neural Networks and Fuzzy Logic — NAFIPS/IFIS/ NASA 94*, San Antonio, TX (pp. 115-120).
- Ponnuswamy, S., Amin, M. B., Jha, R., & Castañón, D. A. (1997). A C3I parallel benchmark based on genetic algorithms implementation and performance analysis. *Journal of Parallel and Distributed Computing*, 47(1), 23-38.
- Sanderson, P. M. (1998, November 29-December 4). Cognitive work analysis and the analysis, design, and evaluation of human computer interactive systems. *Proceedings of the Annual Conference of the Computer-Human Interaction Special Interest Group (CHISIG) of the Ergonomics Society of Australia (OzCHI98)*, Adelaide, South Australia (pp. 40-45).
- Setiono, R. (2000). Extracting M-of-N rules from trained neural networks. *IEEE Transactions on Neural Networks*, 11(2), 512-519.
- Setiono, R., Leow, W. K., & Zurada, J. M. (2002). Extraction of rules from artificial neural networks for nonlinear regression. *IEEE Transactions on Neural Networks*, 13(3), 564-577.
- Steinberg, D., & Colla, P. L. (1995). *CART: Tree-structured non-parametric data analysis*. San Diego, CA: Salford Systems.
- Sugeno, M. (1985). *Industrial applications of fuzzy control*. Amsterdam: Elsevier Science Publishing Company.
- Takagi, T., & Sugeno, M. (1983, December 15-18). Derivation of fuzzy control rules from human operator's control actions. *Proceedings of the IFAC Symposium on Fuzzy Information, Knowledge Representation and Decision Analysis*, Marseilles, France (pp. 55-60).
- Tan, K. C., & Li, Y. (2001). Performance-based control system design automation via evolutionary computing. *Engineering Applications of Artificial Intelligence*, 14(4), 473-486.
- Tan, K. C., Yu, Q., Heng, C. M., & Lee, T. H. (2003). Evolutionary computing for knowledge discovery in medical diagnosis. *Artificial Intelligence in Medicine*, 27(2), 129-154.
- Tran, C., Abraham, A., & Jain, L. (2004). Modeling decision support systems using hybrid neurocomputing. *Neurocomputing*, 61C, 85-97.
- Tran, C., Jain, L., & Abraham, A. (2002a, December 2-6). Adaptation of Mamdani fuzzy inference system using neuro — genetic approach for tactical air combat decision support system. *Proceedings of the 15th Australian Joint Conference on Artificial Intelligence (AI'02)*, Canberra, Australia (pp. 402-410). Berlin: Springer Verlag.
- Tran C., Jain, L., & Abraham, A. (2002b), Adaptive database learning in decision support system using evolutionary fuzzy systems: A generic framework, hybrid information systems. In A. Abraham & M. Oppen (Eds.), *Advances in soft computing* (pp. 237-252). Berlin: Physica Verlag.
- Tran, C., Jain, L., & Abraham, A. (2002c). TACDSS: Adaptation of a Takagi-Sugeno hybrid neuro-fuzzy system. *Proceedings of the 7th Online World Conference on*

- Soft Computing in Industrial Applications (WSC7)* (pp. 53-62). Berlin: Springer Verlag.
- Tran, C., & Jain, L. (2000, September). Intelligent techniques for decision support in tactical environment. *Proceedings of the Warfighting in Complex Terrain: Land Warfare Conference 2000*, Melbourne, Australia, Australian Defence Science & Technology Organisation (pp. 403-411).
- Wang, L. X., & Mendel J. M. (1992). Generating fuzzy rules by learning from examples. *IEEE Transactions on System, Man and Cybernetics*, 22(6), 1414-1427.
- Watts, M., Woodford, B., & Kasabov, N. (1999, November 22-24). FuzzyCOPE: A software environment for building intelligent systems — The past, the present, and the future. *Proceedings of the ICONIP/ANZIIS/ANNES'99 Workshop*, Dunedin/Queenstown, New Zealand (pp. 188-192).
- Zadeh, L. A. (1973). Outline of a new approach to the analysis of complex systems and decision processes. *IEEE Transactions on System, Man, & Cybernetics*, 3(1), 28-44.
- Zadeh, L. A. (1998, August 21-31). Roles of soft computing and fuzzy logic in the conception, design, and deployment of information/intelligent systems. In O. Kaynak, L. A. Zadeh, B. Turksen, & I. J. Rudas (Eds.), *Computational intelligence: Soft computing and fuzzy-neuro integration and applications. Proceedings of NATO Advanced Study Institute on Soft Computing and its Applications: Vol. 162*, Manavgat, Antalya, Turkey (pp. 1-9).
- Zurada, J. M. (1992). *Introduction to artificial neural systems*. St. Paul, MN: West Publishing Company.

Chapter II

Application of Text Mining Methodologies to Health Insurance Schedules

Ah Chung Tsoi, Monash University, Australia

Phuong Kim To, Tedis P/L, Australia

Markus Hagenbuchner, University of Wollongong, Australia

ABSTRACT

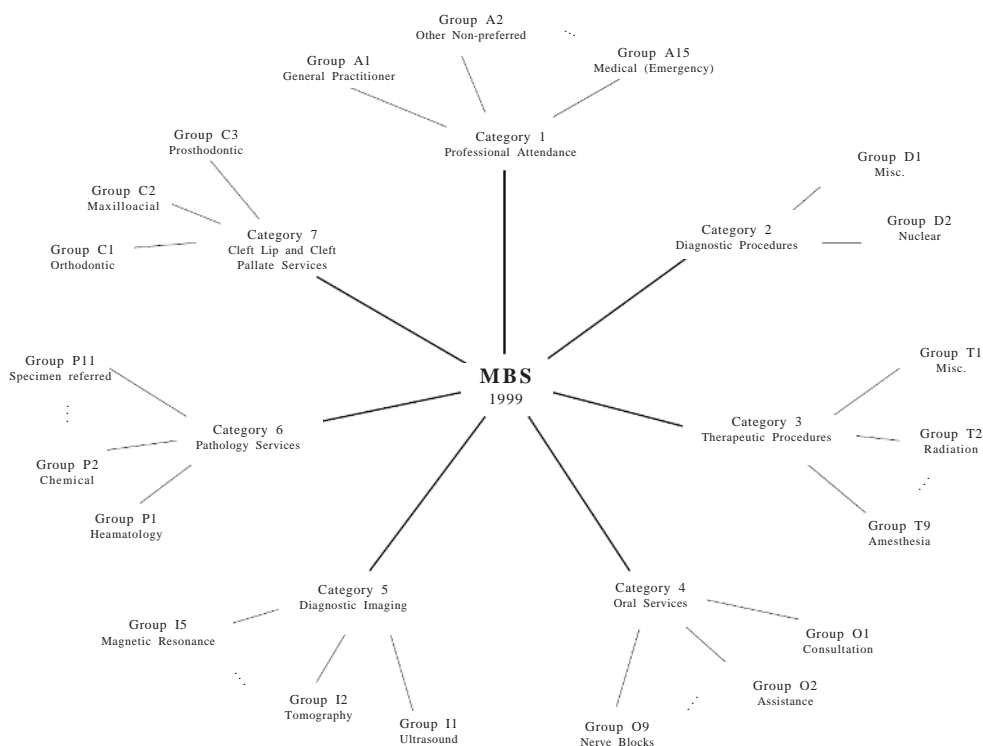
This chapter describes the application of a number of text mining techniques to discover patterns in the health insurance schedule with an aim to uncover any inconsistency or ambiguity in the schedule. In particular, we will apply first a simple “bag of words” technique to study the text data, and to evaluate the hypothesis: Is there any inconsistency in the text description of the medical procedures used? It is found that the hypothesis is not valid, and hence the investigation is continued on how best to cluster the text. This work would have significance to health insurers to assist them to differentiate descriptions of the medical procedures. Secondly, it would also assist the health insurer to describe medical procedures in an unambiguous manner.

AUSTRALIAN HEALTH INSURANCE SYSTEM

In Australia, there is a universal health insurance system for her citizens and permanent residents. This publicly-funded health insurance scheme is administered by a federal government department called the Health Insurance Commission (HIC). In addition, the Australian Department of Health and Ageing (DoHA), after consultation with the medical fraternity, publishes a manual called Medicare Benefit Schedule (MBS) in which it details each medical treatment procedure and its associated rebate to the medical service providers who provide such services. When a patient visits a medical service provider, the HIC will refund or pay the medical service provider at the rate published in the MBS¹ (the MBS is publicly available online from <http://www.health.gov.au/pubs/mbs/mbs/css/index.htm>).

Therefore, the description of medical treatment procedures in the MBS should be clear and unambiguous to interpretation by a reasonable medical service provider as ambiguities would lead to the wrong medical treatment procedure being used to invoice the patient or the HIC. However, the MBS has developed over the years, and is derived through extensive consultations with medical service providers over a lengthy period. Consequently, there may exist inconsistencies or ambiguities within the schedule. In this chapter, we propose to use text mining methodologies to discover if there are any ambiguities in the MBS.

Figure 1. An overview of the MBS structure in the year of 1999



The MBS is divided into seven categories, each of which describes a collection of treatments related to a particular type, such as diagnostic treatments, therapeutic treatments, oral treatments, and so on. Each category is further divided into groups. For example, in category 1, there are 15 groups, A_1, A_2, \dots, A_{15} . Within each group, there are a number of medical procedures which are denoted by unique item numbers. In other words, the MBS is arranged in a hierarchical tree manner, designed so that it is easy for medical service providers to find appropriate items which represent the medical procedures provided to the patient.² This underlying MBS structure is outlined in Figure 1.

This chapter evaluates the following:

- **Hypothesis** — Given the arrangement of the items in the way they are organised in the MBS (Figure 1), are there any ambiguities within this classification? Here, ambiguity is measured in terms of a confusion table comparing the classification given by the application of text mining techniques and the classification given in the MBS. Ideally, if the items are arranged without any ambiguities at all (as measured by text mining techniques), the confusion table should be diagonal with zero off diagonal terms.
- **Optimal grouping** — Assuming that the classification given in MBS is ambiguous (as revealed in our subsequent investigation of the hypothesis), what is the “optimal” arrangement of the item descriptions using text mining techniques (here “optimal” is measured with respect to text mining techniques)? In other words, we wish to find an “optimal” grouping of the item descriptions together such that there will be a minimum of misclassifications.

The benefits of this work are as follows:

- From the DoHA point of view, it will allow the discovery of any existing ambiguities in the MBS. In order to make procedures described in the MBS as distinct as possible, the described methodology can be employed in evaluating the hypothesis in designing the MBS such that there would not be any ambiguities from a text mining point of view. This will lead to a better description of the procedures so that there will be little misinterpretation by medical service providers.
- From a service provider’s point of view, the removal of ambiguities would allow efficient computer-assisted searching. This will limit misinterpretation, and allow the implementation of a semi-automatic process for the generation of claims and receipts.
- While the “optimal grouping” process is mainly derived from a curiosity point of view, this may assist the HIC in re-grouping some of their existing descriptions of items in the MBS, so that there will be less opportunities for misinterpretation.

Obviously, the validity of the described method lies in the validity of text mining techniques in unambiguously classifying a set of documents. Unfortunately, this may not be the case, as new text mining techniques are constantly being developed.

However, the value of the work presented in this paper lies in the ability to use existing text mining techniques and to discover, as far as possible, any ambiguities within the MBS. This is bound to be a conservative measure, as we can only discover ambiguities as far as possible given the existing tools. There will be other ambiguities

which remain uncovered by current text mining techniques. But at least, using our approach will clear up some of the existing ambiguities. In other words, the text mining techniques do not claim to be exhaustive. Instead, they will indicate ambiguities as far as possible, given their limitations.

The structure of this chapter is as follows: In the next section, we describe what text mining is, and how our proposed techniques fall into the general fabric of text mining research. In the following section, we will describe the “bag of words” approach to text mining. This is the simplest method in that it does not take any cognizance of semantics among the words; each word is treated in isolation. In addition, this will give an answer to the hypothesis as stated above. If ambiguities are discovered by using such a simple text mining technique, then there must exist ambiguities in the set of documents describing the medical procedures. This will give us a repository of results to compare with those when we use other text mining techniques. In the next section, we describe briefly the latent semantic kernel (LSK) technique to pre-process the feature vectors representing the text. In this technique, the intention is that it is possible to manipulate the original feature vectors representing the documents and to shorten them so that they can better represent the “hidden” message in the documents. We show results which do not assume the categories as given in the MBS.

TEXT MINING

In text mining, there are two main issues: retrieval and classification (Berry, 2004).

- **Retrieval techniques** — used to retrieve the particular document:
 - *Keyword-based search* — this is the simplest method in that it will retrieve a document or documents which matches a particular set of key words provided by the user. This is often called “queries”.
 - *Vector space-based retrieval method* — this is often called a “bag of words” approach. It represents the document in terms of a set of feature vectors. Then, the vectors can be manipulated so as to show patterns, for example, by grouping similar vectors into clusters (Nigam, McCallum, Thrun, & Mitchell, 2000; Salton, 1983).
 - *Latent semantic analysis* — this is to study the latent or hidden structure of the set of documents with respect to “semantics”. Here “semantics” is taken to mean “correlation” within the set of documents; it does not mean that the technique will discover the “semantic” relationships between words in the sense of linguistics (Salton, 1983).
 - *Probabilistic latent semantic analysis* — this is to consider the correlation within the set of documents within a probabilistic setting (Hofmann, 1999a).
- **Classification techniques** — used to assign data to classes.
 - *Manual classification* — a set of documents is classified manually into a set of classes or sub-classes.
 - *Rule-based classification* — a set of rules as determined by experts is used to classify a set of documents.
 - *Naïve Bayes classification* — this uses Bayes’ theorem to classify a set of documents, with some additional assumptions (Duda, 2001).

- *Probabilistic latent semantic analysis classification* — this uses the probabilistic latent semantic analysis technique to classify the set of documents (Hofmann, 1999b).
- *Support vector machine classification* — this is to use support vector machine techniques to classify the set of documents (Scholkopf, Burges, & Smola, 1999).

This chapter explores the “bag of words” technique to classify the set of documents into clusters and compare them with those given in the MBS. The chapter also employs the latent semantic kernel technique, a technique from kernel machine methods (based on support vector machine techniques) to manipulate the features of the set of documents before subjecting them to clustering techniques.

BAG OF WORDS

If we are given a set of m documents $D = [d_1, d_2, \dots, d_m]$, it is quite natural to represent them in terms of vector space representation. From this set of documents it is simple to find out the set of vocabularies used. In order that the set of vocabularies would be meaningful, care is taken by using the stemmisation technique which regards words of the same stem to be one word. For example, the words “representation” and “represent” are considered as one word, rather than two distinct words, as they have the same stem. Secondly, in order that the set of vocabularies would be useful to distinguish documents, we eliminate common words, like “the”, “a”, and “is” from the set of vocabularies. Thus, after these two steps, it is possible to have a set of vocabularies w_1, w_2, \dots, w_n which represents the words used in the set of documents D . Then, each document can be represented as an n -vector with elements which denote the frequency of occurrence of the word in the document d_i , and 0 if the word does not occur in the document d_i . Thus, from a representation point of view, the set of documents D can be equivalently represented by a set of vectors $V = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m]$, where \mathbf{v}_i is an n -vector. Note that this set of vectors V may be sparse, as not every word in the vocabulary occurs in the document (Nigam et al., 2000). The set of vectors V can be clustered together to form clusters using standard techniques (Duda, 2001).

Table 1. An overview over the seven categories in the MBS

Category	Number of items
1	158
2	108
3	2734
4	162
5	504
6	302
7	62
Total	4030

Table 2. A confusion table showing the classification of documents (the actual classifications as indicated in the MBS are given horizontally; classifications as obtained by the naïve Bayes method are presented vertically)

Category	1	2	3	4	5	6	7	Total	% Accuracy
1	79	0	0	0	0	0	0	79	100.00
2	1	25	9	0	12	7	0	54	46.30
3	12	3	1323	15	10	3	1	1367	96.78
4	1	0	62	18	0	0	0	81	22.22
5	0	3	18	0	229	1	1	252	90.87
6	0	2	0	0	1	148	0	151	98.01
7	3	0	1	1	2	0	24	31	77.42

In our case, we consider each description of an MBS item as a document. We have a total of 4030 documents; each document may be of varying length, dependent on the description of the particular medical procedure. Table 1 gives a summary of the number of documents in each category.

After taking out commonly occurring words, words with the same stem count, and so on, we find that there are a total of 4569 distinct words in the vocabulary.

We will use 50% of the total number of items as the training data set, while the other 50% will be used as a testing data set to evaluate the generalisability of the techniques used. In other words, we have 2015 documents in the training data set, and 2015 in the testing data set. The content of the training data set is obtained by randomly choosing items from a particular group so as to ensure that the training data set is sufficiently rich and representative of the underlying data set.

Once we represent the set of data in this manner, we can then cluster them together using a simple clustering technique, such as the naïve Bayes classification method (Duda, 2001). The results of this clustering are shown in Table 2.

The percentage accuracy is, on average, 91.61%, with 1846 documents out of 2015 correctly classified. It is further noted that some of the categories are badly classified, for example, category-2 and category-4. Indeed, it is found that 62 out of 81 category-4 items are misclassified as category-3. Similarly, 12 out of 54 category-2 items are misclassified as category-5 items.

This result indicates that the hypothesis is not valid; there *are* ambiguities in the description of the items in each category, apart from category-1, which could be confused with those in other categories. In particular, there is a high risk of confusing those items in category-4 with those in category-3.

A close examination of the list of the 62 category-4 items which are misclassified as category-3 items by the naïve Bayes classification method indicates that they are indeed very similar to those in category-3. For simplicity, when we say items in category-3, we mean that those items are also correctly classified into category-3 by the classification method. Tables 3 and 4 give an illustration of the misclassified items. It is noted that misclassified items 52000, 52003, 52006, and 52009 in Table 3 are very similar to the category-3 items listed in Table 4.

Table 3. Category-4 Items 52000, 52003, 52006, and 52009 misclassified by the naïve Bayes method as Category-3 items

Item No	Item Description
52000	Skin and subcutaneous tissue or mucous membrane, repair of recent wound of, on face or neck, small (not more than 7 cm long), superficial
52003	Skin and subcutaneous tissue or mucous membrane, repair of recent wound of, on face or neck, small (not more than 7 cm long), involving deeper tissue
52006	Skin and subcutaneous tissue or mucous membrane, repair of recent wound of, on face or neck, large (more than 7 cm long), superficial
52009	Skin and subcutaneous tissue or mucous membrane, repair of recent wound of, on face or neck, large (more than 7 cm long), involving deeper tissue

Table 4. Some items in Category 3 which are similar to items 52000, 52003, 52006, and 52009

Item No	Item Description
30026	Skin and subcutaneous tissue or mucous membrane, repair of wound of, other than wound closure at time of surgery, not on face or neck, small (not more than 7cm long), superficial, not being a service to which another item in Group T4 applies
30035	Skin and subcutaneous tissue or mucous membrane, repair of wound of, other than wound closure at time of surgery, on face or neck, small (not more than 7cm long), involving deeper tissue
30038	Skin and subcutaneous tissue or mucous membrane, repair of wound of, other than wound closure at time of surgery, not on face or neck, large (more than 7cm long), superficial, not being a service to which another item in Group T4 applies
30041	Skin and subcutaneous tissue or mucous membrane, repair of wound of, other than wound closure at time of surgery, not on face or neck, large (more than 7cm long), involving deeper tissue, not being a service to which another item in Group T4 applies
30045	Skin and subcutaneous tissue or mucous membrane, repair of wound of, other than wound closure at time of surgery, on face or neck, large (more than 7cm long), superficial
30048	Skin and subcutaneous tissue or mucous membrane, repair of wound of, other than wound closure at time of surgery, on face or neck, large (more than 7cm long), involving deeper tissue

It is observed that the way items 5200X are described is very similar to those represented in items 300YY. For example, item 52000 describes a medical procedure to repair small superficial cuts on the face or neck. On the other hand, item 30026 describes the same medical procedure except that it indicates that the wounds are *not* on the face

or neck, with the distinguishing feature that this is not a service to which another item in Group T4 applies. It is noted that the description of item 30026 uses the word “not” to distinguish this from that of item 52000, as well as appending an extra phrase “not being a service to which another item in Group T4 applies”. From a vector space point of view, the vector representing item 52000 is very close³ to item 30026, closer than other items in category-4, due to the few extra distinguishing words between the two. Hence, item 52000 is classified as “one” in category-3, instead of “one” in category-4. Similar observations can be made for other items shown in Table 3, when compared to those shown in Table 4.

Table 5. Some correctly classified Category-1 items

Item No	Item description
3	Professional attendance at consulting rooms (not being a service to which any other item applies) by a general practitioner for an obvious problem characterised by the straightforward nature of the task that requires a short patient history and, if required, limited examination and management -- each attendance
4	Professional attendance, other than a service to which any other item applies, and not being an attendance at consulting rooms, an institution, a hospital, or a nursing home by a general practitioner for an obvious problem characterised by the straightforward nature of the task that requires a short patient history and, if required, limited examination and management -- an attendance on 1 or more patients on 1 occasion -- each patient
13	Professional attendance at an institution (not being a service to which any other item applies) by a general practitioner for an obvious problem characterised by the straightforward nature of the task that requires a short patient history and, if required, limited examination and management -- an attendance on 1 or more patients at 1 institution on 1 occasion -- each patient
19	Professional attendance at a hospital (not being a service to which any other item applies) by a general practitioner for an obvious problem characterised by the straightforward nature of the task that requires a short patient history and, if required, limited examination and management -- an attendance on 1 or more patients at 1 hospital on 1 occasion -- each patient
20	Professional attendance (not being a service to which any other item applies) at a nursing home including aged persons' accommodation attached to a nursing home or aged persons' accommodation situated within a complex that includes a nursing home (other than a professional attendance at a self contained unit) or professional attendance at consulting rooms situated within such a complex where the patient is accommodated in a nursing home or aged persons' accommodation (not being accommodation in a self contained unit) by a general practitioner for an obvious problem characterised by the straightforward nature of the task that requires a short patient history and, if required, limited examination and management -- an attendance on 1 or more patients at 1 nursing home on 1 occasion -- each patient

Table 6. Some correctly classified Category-5 items

Item No	Item description
55028	Head, ultrasound scan of, performed by, or on behalf of, a medical practitioner where: (a) the patient is referred by a medical practitioner for ultrasonic examination not being a service associated with a service to which an item in Subgroups 2 or 3 of this Group applies; and (b) the referring medical practitioner is not a member of a group of practitioners of which the first mentioned practitioner is a member (R)
55029	Head, ultrasound scan of, where the patient is not referred by a medical practitioner, not being a service associated with a service to which an item in Subgroups 2 or 3 of this Group applies (NR)
55030	Orbital contents, ultrasound scan of, performed by, or on behalf of, a medical practitioner where: (a) the patient is referred by a medical practitioner for ultrasonic examination not being a service associated with a service to which an item in Subgroups 2 or 3 of this Group applies; and (b) the referring medical practitioner is not a member of a group of practitioners of which the first mentioned practitioner is a member (R)
55031	Orbital contents, ultrasound scan of, where the patient is not referred by a medical practitioner, not being a service associated with a service to which an item in Subgroups 2 or 3 of this Group applies (NR)
55033	Neck, 1 or more structures of, ultrasound scan of, where the patient is not referred by a medical practitioner, not being a service associated with a service to which an item in Subgroups 2 or 3 of this Group applies (NR)

On the other hand, Tables 5 and 6 show items which are correctly classified in category-1 and category-5 respectively. It is observed that items shown in Table 5 are distinct from those shown in Table 6 in their descriptions. A careful examination of correctly-classified category-1 items, together with a comparison of their descriptions with those correctly-classified category-5 items confirms the observations shown in Tables 5 and 6. In other words, the vectors representing correctly-classified category-1 items are closer to other vectors in the same category than other vectors representing other categories.

SUPPORT VECTOR MACHINE AND KERNEL MACHINE METHODOLOGIES

In this section, we will briefly describe the support vector machine and the kernel machine techniques.

Support Vector Machine and Kernel Machine Methodology

In recent years, there has been increasing interest in a method called support vector machines (Cristianni & Shawe-Taylor, 2000; Guermeur, 2002; Joachims, 1999; Vapnik, 1995). In brief, this can be explained quite easily as follows: Assume a set of (n -

dimensional) vectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$. Assuming that this set of vectors is drawn from two classes, 1 and -1. If these classes are linearly separable, then there exists a straight line dividing these two classes as shown on the left of Figure 2. In Figure 2, it is observed that the vectors are well separated. Now if the two classes cannot be separated by a straight line, the situation becomes more interesting. Traditionally, in this case we use a non-linear classifier to separate the classes as shown on the right of Figure 2. In general terms, any two collections of n -dimensional vectors are said to be linearly separable if there exists an $(n-1)$ -dimensional hyper-plane that separates the two collections.

Figure 2. Illustration of the linear separability of classes (the two classes at top are separable by a single line, as indicated; for the lower two classes there is no line that can separate them)

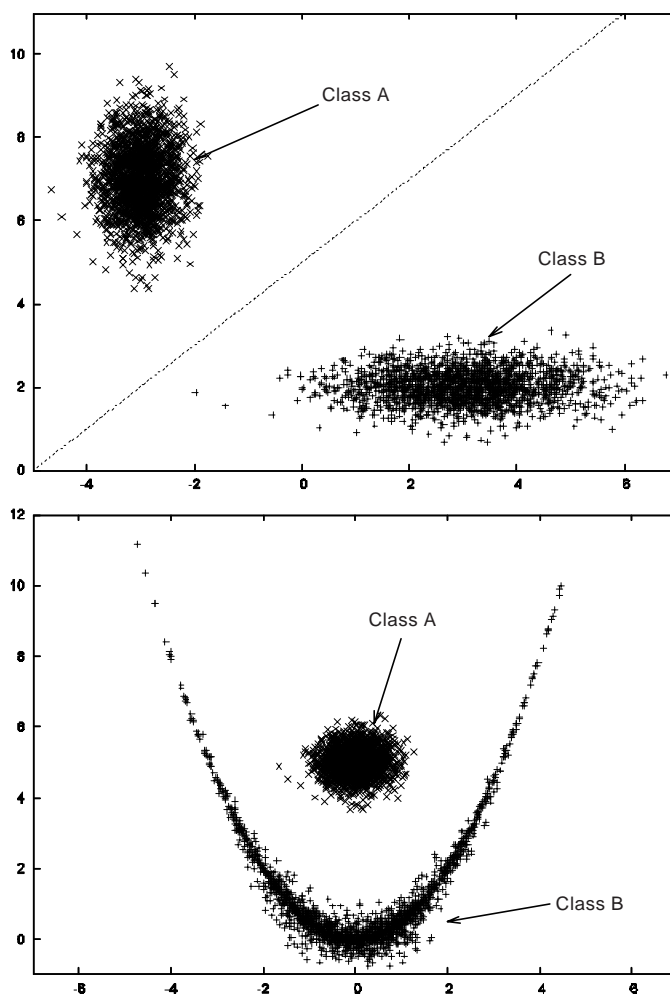


Table 7. Exclusive-OR example

x	y	class
0	0	1
1	1	1
0	1	0
1	0	0

Table 8. Extended exclusive-OR example

x	y	z	class
0	0	0	1
1	1	1	1
0	1	0	0
1	0	0	0

One intuition is inspired by the following example: In the exclusive-OR case, we know that it is not possible to separate the two classes using a straight line, when the problem is represented in two dimensions. However, we know that if we increase the dimension of the exclusive-OR example by one, then in three dimensions one can find a hyper-plane which will separate the two classes. This can be observed in Tables 7 and 8, respectively.

Here it is observed that the two classes are easily separated when we simply add one extra dimension. The support vector machine uses this insight, namely, in the case when it is not possible to separate the two classes by a hyper-plane; if we augment the dimension of the problem sufficiently, it is possible to separate the two classes by a hyper-plane. $f(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x}) + b$, where \mathbf{w} is a set of weights, and b a constant in this high-dimensional space. The embedding of the vectors \mathbf{x} in the high-dimensional plane is to transform them equivalently to $\phi(\mathbf{x})$, where $\phi(\cdot)$ is a coordinate transformation. The question then becomes: how to find such a transformation $\phi(\cdot)$?

Let us define a kernel function as follows:

$$K(\mathbf{x}, \mathbf{z}) = \phi(\mathbf{x})^T \phi(\mathbf{z}) \quad (1)$$

where ϕ is a mapping from X to an inner product feature space F . It is noted that the kernel thus defined is symmetric, in other words $K(\mathbf{x}, \mathbf{z}) = K(\mathbf{z}, \mathbf{x})$. Now let us define the matrix $\mathbf{X} = [\mathbf{x}_1 \mathbf{x}_2 \dots \mathbf{x}_n]$. It is possible to define the symmetric matrix:

$$\mathbf{X}^T \mathbf{X} = \begin{bmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \vdots \\ \mathbf{x}_n^T \end{bmatrix} \begin{bmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \dots & \mathbf{x}_n \end{bmatrix} \quad (2)$$

In a similar manner, it is possible to define the kernel matrix:

$$\mathbf{K} = [\phi(\mathbf{x}_1) \phi(\mathbf{x}_2) \dots \phi(\mathbf{x}_n)]^T [\phi(\mathbf{x}_1) \phi(\mathbf{x}_2) \dots \phi(\mathbf{x}_n)] \quad (3)$$

Note that the kernel matrix \mathbf{K} is symmetric. Hence, it is possible to find an orthogonal matrix \mathbf{V} such that $\mathbf{K} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$, where $\mathbf{\Lambda}$ is a diagonal matrix containing the eigenvalues of \mathbf{K} . It is convenient to sort the diagonal values of $\mathbf{\Lambda}$ such that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. It turns out that one necessary requirement of the matrix \mathbf{K} to be a kernel function is that the eigenvalue matrix $\mathbf{\Lambda}$ must contain all positive entries, in other words, $\lambda_i \geq 0$. This implies that in general, for the transformation $\phi(\cdot)$ to be a valid transformation, it must satisfy some conditions such that the kernel function formed is symmetric. This is known as the Mercer conditions (Cristianni & Shawe-Taylor, 2000).

There are many possible such transformations; some common ones (Cristianni & Shawe-Taylor, 2000) being:

Power kernel: $K(\mathbf{x}, \mathbf{z}) = (K(\mathbf{x}, \mathbf{z}) + c)^p$ where $p = 2, 4, \dots$

Gaussian kernel: $K(\mathbf{x}, \mathbf{z}) = \exp\left(\frac{K(\mathbf{x}, \mathbf{x}) + K(\mathbf{z}, \mathbf{z}) - 2K(\mathbf{x}, \mathbf{z})}{\sigma^2}\right)$.

There exist quite efficient algorithms using optimisation theory which will obtain a set of support vectors and the corresponding weights of the hyper-plane for a particular problem (Cristianni & Shawe-Taylor, 2000; Joachims, 1999). This is based on re-formulating the problem as a quadratic programming problem with linear constraints. Once it is thus re-formulated, the solutions can be obtained very efficiently.

It was also discovered that the idea of a kernel is quite general (Scholkopf, Burges, & Smola, 1999). Indeed, instead of working with the original vectors \mathbf{x} , it is possible to work with the transformed vectors $\phi(\mathbf{x})$ in the feature space, and most classic algorithms, for example, principal component analysis, canonical correlation analysis, and Fisher's discriminant analysis, all have equivalent algorithms in the kernel space. The advantage of working in the feature space is that the dimension is normally much lower than the original space.

Latent Semantic Kernel Technique

The latent semantic kernel method follows the same trend as the kernel machine methodology (Cristianini, Lodhi, & Shawe-Taylor, 2002). The latent semantic kernel is the kernel machine counterpart of the latent semantic technique, except that it operates in a lower dimension feature space, and hence is more efficient.

In latent semantic analysis, we have a set of documents represented by D , in terms of $V = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m]$. This set of vectors can be concatenated into a matrix \mathbf{D} , then we may apply a singular value decomposition on the matrix \mathbf{D} as follows:

$$\mathbf{D} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T \quad (4)$$

where \mathbf{D} is a $n \times m$ matrix, \mathbf{U} is an ortho normal $n \times n$ matrix such that $\mathbf{U}\mathbf{U}^T = \mathbf{I}$, \mathbf{V} is an ortho normal $m \times m$ matrix, such that $\mathbf{V}^T\mathbf{V} = \mathbf{I}$, and $\mathbf{\Sigma}$ is a $n \times m$ matrix, with diagonal entries $\sigma_1, \sigma_2, \dots, \sigma_n$, if $n > m$ or $\sigma_1, \sigma_2, \dots, \sigma_m$ if $m > n$.

Often the singular values σ are arranged so that $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$. Thus, the singular values σ give some information on the “energy” of each dimension. It is possible that some of the σ may be small or negligible. In this case, it is possible to say that there are only a few significant singular values. For example, if we have $\sigma_1 \geq \sigma_2 \geq \sigma_i \gg \sigma_{i+1} \geq \sigma_n$, then it is possible to approximate the set by $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_i$. In this case, it is possible to perform a dimension reduction on the original data so that it conforms to the reduced order data. By reducing the order of representation, we are “compressing” the data, thus inducing it to have a “semantic” representation.

The idea behind the latent semantic kernel is, instead of considering the original document matrix \mathbf{D} , to consider the kernel matrix $\mathbf{K} = \mathbf{D}^T\mathbf{D}$. In this case, the feature space dimension is smaller, as we normally assume that there are more words in the vocabulary than the number of documents, in other words, $n \gg m$. Thus, by operating on $m \times m$ matrix $\mathbf{D}^T\mathbf{D}$, it is a smaller space than the original n dimensional space. Once it is recognised that the kernel is K , we can then operate on this kernel, for example, performing singular value decomposition on the matrix \mathbf{K} , and find the corresponding singular values.

One particular aspect of performing a singular value decomposition is to find a reduced order model such that, in the reduced order space it will approximate the model in the original space in the sense that the approximated model contains most of the “energies” in the original model. This concept can also be applied to the latent semantic kernel technique. It is possible to find a reduced order representation of the original features in such a manner that the reduced order representation contains most of the “energies” in the original representation.

The latent semantic kernel algorithm (Cristianini, Lodhi, & Shawe-Taylor, 2002) can be described as follows:

Given a kernel K , training set d_1, \dots, d_m and a number T :

for $i = 1$ to m do:

$$norm2[i] = K(d_i, d_i);$$

for $j = 1$ to T do:

$$i_j = \arg \max_i (norm2[i]);$$

$$index[j] = i_j;$$

$$size[j] = \sqrt{norm2[i_j]};$$

for $i = 1$ to m do:

$$feat[i, j] = (K(d_i, d_{i_j}) - \sum_{t=1}^{j-1} feat[i, t] * feat[i_j, t]) / size[j];$$

```

    norm2[i] = norm2[i] - feat(i,j)* feat(i,j);

    end;

end;

return feat[i, j] as the  $j^{th}$  feature of input  $i$ ;

To classify a new example  $d$ :

    for  $j = 1$  to  $T$  do:

        newfeat[j] =  $(K(d, d_{i_j}) - \sum_{t=1}^{j-1} newfeat[t] * feat[i_j, t]) / size[j]$ ;

    end;

return newfeat[j] as the  $j^{th}$  feature of the example  $d$ .

```

In our work, we use the latent semantic kernel method as a pre-processing technique in that we have a set of documents represented by the matrix \mathbf{D} , and as mentioned previously, the matrix is quite sparse in that there are many null values within the matrix. Hence, our aim is to project this representation onto a representation in a reduced order space so that most of the “energies” are still retained. Once we obtain this set of reduced order representation, we can then manipulate them in the same way as we manipulate the full feature vectors.

Application of Latent Semantic Kernel Methodology to the Medical Benefit Schedule

We wish to apply the latent semantic kernel technique to the set of descriptions of items as contained in the MBS. The vector representation as described in in the *Bag of Words* section is used in the latent semantic kernel (LSK) approach, and it is known as the full feature vectors or, in short, full features.

We apply the LSK algorithm to the *full features* and produce reduced features which have a dimension determined by a variable T ($T \leq n$, the number of words in the corpus). In this paper, we will use the term LSK reduced features to represent such features.

We use the following procedures in our experiments with the latent semantic kernel:

- Run **bag of words method** to produce full features.
- Run **LSK algorithm** to produce LSK reduced features.
- Experiments with both **LSK** and **full features** including:
 - *Binary classification* — this will allow only binary classification using support vector machine techniques
 - *Multi-classification* — this will allow multi-class classification using support vector machine techniques
 - *Clustering the items in the MBS* using both full features and reduced features
 - *Compare the clustering result* using multi-classification results

EXPERIMENTS USING BINARY CLASSIFICATION

In this section we report the results of using a binary classification support vector machine technique for the classification of the items (Joachims, 1999). This is interesting in that it shows us the results of assuming one class, and the other items are assumed to be in a different class.

Originally, items in the MBS are classified into seven categories: 1, 2,..., 6 and 7. We have trained a binary classifier for both the full features and for the reduced features regarding each category versus the others. For example, we use category-1 and assume all the other categories are grouped as another class.

We run experiments on reduced features for each category where the dimension T of the reduced features was chosen from a set of values within the range [16; 2048].

We show the results of category-1 versus other categories first, and then summarise other situations and draw some general conclusions concerning the experiments.

For the experiment about category-1 versus others (see Table 9), we observe that the accuracy climbs from $T = 16$ rapidly to a stable position of a maximum at $T = 32$. Note that even though the accuracies oscillate about this maximum value for later values of T , this can be observed to be minor, and can be attributed to noise. From this, we can conclude that if we classify category-1 versus all the other categories, then using $T = 32$ is sufficient to capture most of the gain. What this implies is that most of the energies in the document matrix are already captured with 32 reduced order features. This shows that the LSK algorithm is very efficient in compressing the features from 4569 words in the vocabulary to only requiring 32 reduced features. Note that it is not possible to

Table 9. Accuracy on the training data and testing data set for Category 1 with various values of T and full features

T	Train		Train	
	Accuracy	Correct (out of 2015)	Accuracy	Correct (out of 2015)
16	99.60	2007	99.31	2001
32	100	2015	99.75	2010
64	100	2015	99.75	2010
128	100	2015	99.80	2011
256	100	2015	99.80	2011
400	100	2015	99.70	2009
512	100	2015	99.75	2010
800	100	2015	99.75	2010
1024	100	2015	99.75	2010
1200	100	2015	99.80	2011
1400	100	2015	99.75	2010
1600	100	2015	99.75	2010
1800	100	2015	99.75	2010
2000	100	2015	99.75	2010
2048	100	2015	99.75	2010
full features	100	2015	99.75	2010

interpret the reduced features, as by nature they consist of a transformed set of features.

For experiments involving other sets, for example, using category-2 versus all the other categories, and so on, similar conclusions to that shown for category-1 are observed. It is found that in general a small value of T is sufficient to capture most of the energies in the document matrix. Each category versus the rest peaks at a particular value of T . On average it is found that $T = 128$ would capture most of the gain possible using the reduced features.

These experiments show us that the LSK algorithm is a very efficient pre-processing unit. The observed efficiency is likely due to the sparseness of the full feature space. It can capture most of the energies contained in the document matrix using a small reduced feature set, in other words, with a value of $T = 128$.

EXPERIMENTS USING MULTIPLE CLASSIFICATIONS

In this section, we report on experiments with multi-classification using support vector machine (SVM) methodology. We first discuss the generalisation of SVM's to multi-class classification, then we describe the experimental results.

SVM, as a method proposed in Vapnik (1995), is suitable for two-class classifications. There are a number of extensions which extend this method to multi-class classification problems (Crammer & Singer, 2001; Guermeur, 2002; Lee, Lin, & Wahba, 2002). It turns out that it is possible that, instead of weighing the cost function as an equal cost (indicating that both classes are equally weighed in a two-class classification problem), one can modify the cost function and weigh the misclassification cost as well (Lee, Lin & Wahba, 2002). Once formulated in this manner, the usual formulation of SVM's can be applied.

In the *Experiments Using Binary Classification* section, we showed that by using a small reduced feature such as $T = 128$, we can capture most of the energies in the document matrix. The following result is obtained by running multi-class classification using support vector machine on reduced features with the same training and testing data

Table 10. A confusion table showing the classification of documents (the row gives the actual classification as indicated in the MBS, while the column shows figures which are obtained by using the support vector machine)

Category	1	2	3	4	5	6	7	total	% accuracy
1	74	0	0	0	0	4	1	79	93.67
2	0	28	13	0	8	5	0	54	51.85
3	1	6	1310	16	20	2	12	1367	95.83
4	0	0	71	9	0	0	1	81	11.11
5	0	11	25	0	211	4	1	252	83.73
6	0	5	5	1	4	136	0	151	90.07
7	0	0	4	0	0	2	25	31	80.65

sets as previously. The average accuracy is 88.98%, with 1793 correctly classified out of 2015 (Table 10).

Note that once again some of the HIC categories are poorly classified. For example, out of 81 HIC category-4 items, 71 are classified as category-3, while only 9 are classified as category-4. This further confirms the results obtained in the *Bag of Words* section, namely that there are ambiguities in the HIC classifications of item descriptions.

EXPERIMENTS WITH CLUSTERING ALGORITHMS

So far, we have experimented on MBS items which are classified by categories as contained in the MBS. We have shown that the MBS contained ambiguities using the bag of words approach. In this section, we ask a different question: If we ignore the grouping of the items into categories as contained in the MBS, but instead we apply clustering algorithms to the item descriptions, how many clusters would we find? Secondly, how efficient would these clusters be in classifying the item descriptions? Efficiency is measured in terms of the confusion matrix in classifying unseen data.

The methodology which we used is as follows:

- Use a clustering algorithm to cluster the document matrix into clusters, and label the clusters accordingly.
- Evaluate the efficiency of the clustering algorithm using a support vector machine. The efficiency is observed by examining the resulting confusion matrix.

The main reason why we need to take a two-step process is that the clustering algorithm is an unsupervised learning algorithm and that we do not have any *a priori* information concerning which item should fall into which cluster. Hence, it is very difficult to evaluate the efficiency of the clusters produced. In our methodology, we evaluate the efficiency of the clusters from the clustering algorithm by assuming that the clusters formed are “ideal”, label them accordingly, and use the SVM (a supervised training algorithm) to evaluate the clusters formed.

Clustering Using Full Features: Choice of Clustering Method

The first experiment was performed on the full features (in other words, using the original document matrix). Different clustering methods were evaluated using various criteria; we found that the repeated bisection method of clustering gives the best results. Hence, we choose to use this clustering method for all future experiments.

In the clustering algorithm (Karypis, 2003; Zhao & Karypis, 2002), the document matrix is first clustered into two groups. One of the groups is selected and bisected further. This process continues until the desired number of clusters is obtained. During each step, the cluster is bisected so that the resulting two-way clustering solution optimises a particular clustering criterion. At the end of the algorithm, the overall optimisation function is minimised. There are a number of optimising functions which can be used. We use a simple criterion which measures the pair-wise similarities between two documents S_i and S_j as follows:

$$\sum_{d_q \in D_i, d_r \in D_j} \cos(d_q, d_r) = \sum_{d_q \in D_i, d_r \in D_j} d_q^T d_r \quad (5)$$

For a particular cluster S_r of size n_r , the entropy is defined as:

$$E(S_r) = -\frac{1}{\log q} \sum_{i=1}^q \frac{n_r^i}{n_r} \log \frac{n_r^i}{n_r} \quad (6)$$

where q is the number of classes in the data set, and n_r^i is the number of documents of the i -th class that were assigned to the r -th cluster. The entropy of the entire clustering solution is defined as:

$$E = \sum_{r=1}^k \frac{n_r}{n} E(S_r) \quad (7)$$

Perfect clustering means that each cluster will contain only one type of document, that is, each document in the same class belong to the same type. In this case, the entropy will be zero. In general, this is impossible. The clustering result which provides the lowest entropy would be the best clustering result. In this chapter, we made use of the Cluto software (Karypis, 2003) for performing the clustering.

Clustering Using Reduced Features into k-Clusters and Use of Support Vector Machine to Evaluate Cluster Accuracy

For experiments in this section, we used reduced features with $T = 128$, together with the repeated bisection clustering method. Our aim was to use a clustering algorithm with the repeated bisection method to group item descriptions into clusters irrespective of their original classification in the MBS. We divided the item descriptions into k (constant) clusters.

The detailed experimental procedures are as follows:

- Use reduced features of all MBS items obtained using the latent semantic kernel method as inputs into the clustering algorithm in order to group items into k -clusters, where k is a variable determined by the user.
- The output from the clustering algorithm gives k clusters. We perform clustering, do a classification on MBS items using the reduced features and the SVM methodology, however in this case cluster the item belonging to results from the clustering method, not the MBS category.
- The classification output can be displayed in a confusion table which can inform us on how well the clustering algorithm has performed.

In this section, we use the following notations:

- **Cluster category** — this is the cluster obtained using the clustering algorithm.
- **HIC category** — this is the category which is provided by the MBS.

Clusters ($k = 7$)

First, we ran a clustering algorithm to create seven cluster categories. From Table 11, we observe how items in HIC categories are distributed into cluster categories.

From this it is observed that the clusters as obtained by the clustering algorithm are quite different from those indicated in the HIC category.

We then validated the classification results from the clustering algorithm using SVM's. This informed us of the quality of the clustering algorithm in grouping the item descriptions. A classification accuracy of 93.70% on the testing data set (that is, 1888 correct out of 2015) was found.

Note that for these experiments, we used the 50% training and 50% testing data sets as in all previous experiments. The distribution of items in the training and testing data sets for each cluster category was selected using a random sampling scheme within each identified category.

Clusters ($k = 8$)

In this section, we experimented with eight clusters instead of seven. This provided the clustering algorithm with more freedom to choose to cluster the item descriptions. When we chose seven clusters, we implicitly tell the cluster algorithm that *no matter how the underlying data look like*, we nevertheless only allow seven clusters to be found. In this manner, even though the underlying data may be more conveniently clustered into a higher number of clusters, by choosing seven clusters we force the cluster algorithm

Table 11. Distribution of MBS items into cluster categories (the cluster categories are given horizontally, and HIC categories are presented vertically)

Class	1	2	3	4	5	6	7
1	0	3	10	0	230	0	0
2	10	39	442	16	107	69	6
3	0	0	1269	62	9	1	3
4	148	13	65	2	7	17	18
5	0	4	297	34	20	0	12
6	0	12	505	47	18	35	21
7	0	37	146	1	113	180	2

Table 12. A confusion table obtained using the support vector machine method

Class	1	2	3	4	5	6	7	total	% accuracy
1	112	0	1	0	1	0	6	120	93.33
2	0	312	2	4	11	5	4	338	92.31
3	0	3	661	2	8	6	3	683	96.78
4	2	4	0	119	3	2	1	131	90.84
5	0	1	1	5	173	1	4	185	93.51
6	1	1	6	2	9	282	4	305	92.46
7	0	13	0	4	6	1	229	253	90.51

Table 13. Distribution of MBS items into eight cluster categories (the cluster categories are given horizontally, and HIC categories are presented vertically)

Class	1	2	3	4	5	6	7
1	0	4	10	0	244	0	0
2	10	39	326	17	104	70	6
3	0	0	1264	61	8	1	3
4	0	2	223	3	9	1	0
5	148	12	53	2	7	16	18
6	0	4	290	32	20	0	12
7	0	13	501	46	15	30	21
8	0	34	67	1	97	184	2

Table 14. A confusion table using support vector machine to validate the quality of cluster categories

Class	1	2	3	4	5	6	7	8	total	% accuracy
1	127	0	1	0	0	1	0	3	132	96.21
2	0	259	2	0	3	5	1	5	275	94.18
3	0	3	657	1	1	11	3	4	680	96.62
4	0	0	0	106	0	7	0	2	115	92.17
5	2	2	0	2	108	4	5	2	125	86.40
6	0	2	1	3	2	170	2	0	180	94.44
7	0	3	5	4	2	8	280	1	303	92.41
8	0	3	0	4	2	5	1	190	205	92.68

to merge the underlying clusters into seven. On the other hand, if we choose eight clusters, this provides more freedom for the clustering algorithm to cluster the underlying data. If it is truly seven clusters, then the algorithm will report that there are seven clusters found. On the other hand, if the data is more conveniently clustered into eight clusters, then by choosing to allow for the possibility of eight clusters, it will allow the clustering algorithm to find one.

First, we ran a clustering algorithm to create eight cluster categories (Table 13).

From Table 13, it is observed that the clusters as obtained by the clustering algorithm are quite different from those indicated in the HIC category.

We then validated the classification results from the clustering algorithm using SVM's. This informed us of the quality of the clustering algorithm in grouping the item descriptions. A classification accuracy is 94.14% (in other words, 1897 correct out of 2015 resulted).

Summary of Observations

It was observed that:

1. The categories as given by the clustering algorithm are good in grouping the medical item descriptions together. The evaluation using the SVM method shows that it is accurate in grouping them together, as the confusion table has dominant diagonal elements, with few misclassified ones. The SVM is a supervised method. Once it is trained, it can be used to evaluate the generalisation capability of the model on testing results with known classifications. Thus, by examining the confusion table produced, it is possible to evaluate how well the model classifies unseen examples. If the confusion table is diagonally dominant, this implies that there are few misclassifications. On the other hand, if the confusion table is not diagonally dominant, this implies that there are high numbers of misclassifications. In our experiment, we found that the confusion table is diagonally dominant, and hence we can conclude that the grouping as obtained by the clustering algorithm is accurate. In other words, the clustering algorithm was able to cluster the underlying data together into reasonably homogeneous groupings.
2. The classification given by the clustering algorithm increases with the number of clusters, in other words, the degree of freedom in which the clustering algorithm is provided. Obviously there will be an upper limit to a reasonable number of clusters used, beyond which there will not be any further noticeable increase in the classification accuracy. It is observed that the accuracy of assuming seven clusters (93.70%) and the accuracy of using eight clusters (94.14%) are very close to one another. Hence we can conclude that the “optimum” number of clusters is around seven or eight.
3. It is noted that the items in the HIC categories are distributed in the cluster categories. This confirms our finding in the *Bag of Words* section that the HIC categories are not “ideal” categories from a similarity point of view, in that it can induce confusion due to their similarity.

CONCLUSION

We have experimented with classification and clustering on full features and reduced features using the latent semantic kernel algorithm. The results show that the LSK algorithm works well on Health Insurance Commission schedules. It has been demonstrated that the HIC categories are ambiguous in the sense that some item descriptions in one category are close to those of another. This ambiguity may be the cause of misinterpretation of the Medicare Benefits Schedule by medical service providers, leading to wrong charges being sent to the HIC for refund. It is shown that by using clustering algorithms, the item descriptions can be grouped into a number of clusters, and moreover, it is found that seven or eight clusters would be sufficient. It is noted, however, that the item descriptions as grouped using the clustering algorithm are quite different to those of the HIC categories. This implies that if the HIC wishes to re-group item descriptions, it would be beneficial to consider the clusters as grouped by using clustering algorithms.

Note that one may say that our methodology is biased towards the clustering algorithm or classification methods because we only use the simple top HIC categories — categories 1 through 7 as shown in Figure 1. This is only a very coarse classification. In actual fact, the items are classified in the MBS according to a three-tiered hierarchical tree as indicated in Figure 1. For example, an item belongs to category- x group- y item number z , where x ranges from 1 to 7, y ranges from 1 to y_i (where i indicates the category that it is in), and z ranges from 1 to z_j (where j depends on which group the item is located). This is a valid criticism in that the top HIC categories may be too coarse to classify the item descriptions. However, at the current stage of development of text mining methods, it is quite difficult to consider the hierarchical coding of the algorithms, especially when there are insufficient numbers of training samples. In the MBS case, a group may contain only a few items. Thus there is an insufficient number of data to train either the clustering algorithm or SVM. Hence our approach in considering the top HIC categories may be one possible way to detect if there are any inconsistencies in the MBS given the limitations of the methodology.

Even though in this chapter we have concentrated on the medical procedure descriptions of a health insurer, the techniques can be equally applicable to many other situations. For example, our developed methodology can be applied to tax legislation, in identifying whether there are any ambiguities in the description of taxable or tax-exempted items. Other applications of our methodology include: description of the social benefit schedule in a welfare state, and the description of degree rules offered by a university.

ACKNOWLEDGMENT

Financial support is gratefully acknowledged from both the Australian Research Council and the Health Insurance Commission through an ARC Linkage Project grant.

REFERENCES

- Australian Department of Health and Ageing. The Medicare Benefits Schedule Online. Retrieved from <http://www.health.gov.au/pubs/mbs/mbs/css/index.htm>
- Berry, M. W. (Ed.). (2004). *Survey of text mining: Clustering, classification, and retrieval*. New York: Springer Verlag.
- Cristianini, N., & Shawe-Taylor, J. (2000). *Support vector machines*. Cambridge, UK: Cambridge University Press.
- Crammer, K., & Singer, Y. (2001). On the algorithmic implementation of multi-class kernel-based vector machines. *Journal of Machine Learning Research*, 2, 265-292.
- Cristianini, N., Lodhi, H., & Shawe-Taylor, J. (2002). Latent semantic kernels. *Journal of Intelligent Information Systems*, 18(2), March, 127-152.
- Duda, R. O., Hart, P. E., & Stock, D. G. (2001). *Pattern classification*. New York: John Wiley & Sons.
- Guermeur, Y. (2002). *A simple unifying theory of multi-class support vector machines* (Research Rep. No. RR-4669). Lorraine, France: Institut National de Recherche Informatique et en Automatique (INRIA), December.

- Hofmann, T. (1999a). Probabilistic latent semantic analysis. *Proceedings of the 15th Conference on Uncertainties in Artificial Intelligence*. Stockholm, Sweden, July 30-August 1 (pp. 289-296).
- Hofmann, T. (1999b). Probabilistic latent semantic indexing. *Proceedings of the 22nd ACM International Conference on Research and Development in Information Retrieval (SIGIR)*. University of Berkeley, CA, August 15-19 (pp. 50-57).
- Joachims, T. (1999). Making large-scale SVM learning practical. In B. Scholkopf, C. J. C. Burges, & A. J. Smola (Eds.), *Advances in kernel methods*. Boston: MIT Press.
- Karypis, G. (2003). *Cluto, a clustering toolkit* (Tech. Rep. No. 02-017). Minneapolis, MN: Department of Computer Science and Engineering, University of Minnesota, November.
- Lee, Y., Lin, Y., & Wahba, G. (2002). *Multi-category support vector machines, theory and application to the classification of microarray data and satellite radiance data* (Tech. Rep. No. 1064). Madison, WI: Department of Statistics, University of Wisconsin.
- Nigam, K., McCallum, A., Thrun, S., & Mitchell, T. (2000). Text classification from labelled and unlabelled documents using EM. *Machine Learning*. 39, 103-134.
- Salton, G., & McGill, M. J. (1983). *Introduction to modern information retrieval*. New York: McGraw Hill.
- Scholkopf, B., Burges, C. J. C., & Smola, A. J. (1999). *Advances in kernel methods*. Boston: MIT Press.
- Vapnik, V. (1995). *The nature of statistical learning theory*. New York: Springer Verlag.
- Zhao, Y., & Karypis, G. (2002). *Criterion function for document clustering, experiments and analysis* (Tech. Rep. No. 01-40). Minneapolis, MN: Department of Computer Science and Engineering, University of Minnesota, February.

ENDNOTES

- ¹ Note that this scheduled cost is somewhat different from those set by the medical service provider associations. A medical service provider can elect to charge at the scheduled cost set by the medical service provider associations, or a fraction of it, and any gaps in between the charge and the refund by the HIC will need to be met either by the patient, or through additional medical insurance cover specifically designed to cover the gap payment.
- ² Note that the Medical Benefit Schedule is a living document in that the schedules are revised once every few years with additional supplements once every three months. The supplements contain minor modifications to particular sections of the schedule, while the major revisions may contain re-classification of the items, deletion or addition of items, mainly for newly introduced medical services, due to technological advances, or clarification of the intent of existing items. The version of MBS used in this chapter is based on the November, 1998, edition with supplements up to and including June 1, 2000.
- ³ Close here means the cosine of the angle between the two vectors is close to 1.

Chapter III

Coordinating Agent Interactions Under Open Environments

Quan Bai, University of Wollongong, Australia

Minjie Zhang, University of Wollongong, Australia

ABSTRACT

An intelligent agent is a reactive, proactive, autonomous, and social entity. The social ability of an agent is exercised in a multi-agent system (MAS), which constitutes a collection of such agents. Current multi-agent systems mostly work in complex, open, and dynamic environments. In an open environment, many facts, such as domain constraints, agent number, and agent relationships, are not fixed. That brings a lot of difficulties to coordinate agents' interactions and cooperation. One major problem that impedes agent interaction is that most current agent interaction protocols are not very suitable for open environments. In this chapter, we introduce an approach to ameliorate agent interactions from two perspectives. First, the approach can enable agents to form knowledge "rich" interaction protocols by using ontologies. Second, we use coloured Petri net (CPN) based methods to enable agents to form interaction protocols dynamically, which are more suitable for agent interaction under open environments.

INTRODUCTION

It is beyond dispute that multi-agent systems are one of the most important design concepts for today's software. A multi-agent system (MAS) is a computational system that constitutes a collection of intelligent agents. An intelligent agent is a reactive, proactive, autonomous, and social entity, which performs a given task using information gleaned from its environment. In general, intelligent agents possess four major properties (Rao & Georgeff, 1992):

- **Reactivity** — agents can perceive their environment and respond in a timely fashion to changes that occur in it;
- **Pro-activity** — agents not only can simply act in response to their environments, but also are able to exhibit goal-directed behaviours by taking the initiative;
- **Autonomy** — agents have some level of self-control ability, and they can operate without the direct intervention of humans; and
- **Social ability** — agents interact with other agents.

The social ability of an agent is exercised in an MAS. An MAS can be considered as a society of agents that live and work together. In such a multi-agent society, interactions between agents are unavoidable (Lesser, 1999). The interaction between agents occurs when an agent has some intentions and has decided to satisfy these through influencing other agents. Agent interactions are established through exchanging messages that specify the desired performatives of other agents and declarative representations of the contents of messages.

The messages exchanged among agents are composed in agent communication languages (ACLs), such as Knowledge Query and Manipulation Language (KQML) (Finin, Labrou, & Mayfield, 1997) and the Foundation for Intelligent Physical Agents (FIPA) ACL (FIPA, 2004). In addition, messages exchanged between agents need to follow some standard patterns, which are described in agent interaction protocols (Cranefield, Purvis, Nowostawski, & Hwang, 2002). As the application domains of MASs are getting more and more complex, many current agent interaction protocols exhibit some limitations that impede MAS implementations. Firstly, many current application domains of MASs require agents to work in changing and uncertain (open) environments. In such environments, interactions between agents may be influenced by some unexpected factors, such as unexpected messages, loss of messages, or deviation in the message order. Most current agent interaction protocols lack mechanisms to handle these unexpected factors. Secondly, agent architectures in some MASs are heterogeneous, and different agents may possess different interaction protocols. Therefore, due to the heterogeneity, when an agent initialises an interaction with others, it cannot guarantee that its interaction protocol can be understood and accepted by other agents. Thirdly, most agents are hard-coded using interaction protocols, which leads to problems. More specifically, issues such as when to use a particular protocol, what information to transmit, what order to execute tasks, and so on, are left to agent designers. This feature reduces the flexibility of the agent interactions because protocols are hard to modify at runtime once they are pre-coded into the agents. Finally, many current interaction protocols, such as KQML, are not specifically designed to carry knowledge. This kind of knowledge “poor” (Lesser, 1998) protocol is not suitable for applications that need

to exchange complex knowledge. In other words, lack of flexibility and robustness of many current interaction protocols greatly limits the implementation of MASs. Accordingly, how to build a flexible and knowledge “rich” interaction protocol has become one of the main research issues in the area of MASs.

To address some of the above limitations, in this chapter we introduce an approach for agent interactions that can ameliorate agent interactions from two perspectives. First, the approach can enable agents to form knowledge “rich” interaction protocols. Toward this objective, we use ontologies to represent knowledge of agents and ontology facilitators to assist agents to search, acquire, and generate ontologies. Second, we develop a coloured Petri net (CPN) based approach to enable agents to form interaction protocols dynamically, which means protocols are not hard-coded within agents but generated by agents according to their capabilities and status.

The remainder of the chapter is arranged as follows. In the second section, we present the concept of ontologies, the formal expressions of ontologies, and the general framework for ontology-based MASs. Basic descriptions of Petri nets (PNs), CPNs, and how to use CPNs to model agent protocol, are presented in the third section. The CPN-based approach, which enables agents to form flexible protocols dynamically, is introduced in the fourth section. In the fifth section, we discuss the potential applications of methods introduced in this chapter. Finally, the conclusion and future direction of this work are presented in the last section.

MULTI-AGENT SYSTEM ONTOLOGY AND KNOWLEDGE LEVEL AGENT INTERACTIONS

To achieve knowledge-level interactions, agents need shared vocabularies to compose their knowledge and conceptualizations of the domain of interest. These conceptualizations can be expressed in ontologies, which are usually defined formally by way of a semantic web or programming languages. To enable agents to establish knowledge “rich” interaction protocols, it is necessary to define common ontologies of the working domain of the MAS, and include ontology facilitators in MASs to assist agents to search, acquire, and generate ontologies.

MAS Ontology and Formal Ontology Expressions

The term “ontology” is borrowed from philosophy. In the context of MASs, an ontology is defined as a computer-readable description of knowledge about concepts, relationships, and constraints that can exist for an agent or a community of agents. In general, ontologies of an MAS can be classified as common or special ontologies. A common ontology is a description of common knowledge that is related to the MAS working domain or the multi-agent society; a special ontology is a description of some specific knowledge that is related to some particular agents or a single agent of the MAS. During the MAS design phase, it is necessary to define both common ontologies of the system working domain as well as ontology representation format. Benefiting from knowledge engineering, common ontologies of many application domains have already been defined in some ontology libraries, established by several distinguished knowledge engineering research institutes (for example, the Stanford KSL Ontolingua Server).

Therefore, the common ontologies of an MAS domain could be defined by the MAS designer, or by reference to existing ontology libraries.

Ontologies are machine-readable knowledge specifications, and they are usually described by formal languages, such as programming or semantic web languages. Currently, there is no agreed standard for ontology representation (format); however, several exclusive ontology languages such as OIL (Fensel, Harmenlen, Horrocks, McGuinness, & Patel-Schneider, 2001) and DAML+OIL (Horrocks & van Harmelen, 2001), have been widely and successfully applied in many application areas. In Bai and Zhang

Figure 1. An ontology example

```
ontology-container
  title "printer product"
  creator "eAuct"
  description "A example ontology for item
    of the auction site"
  publisher "eAuct"
  type "ontology"
  format "OIL"
  source "http://www.eauct.com.au"
  language "en-au"

ontology-definitions
  class-def defined item
  slot-def itemID
    domain item
  class-def defined digitalProduct
    subclass-of item
  class-def printer
    subclass-of
    computerDigitalProduct
  slot-def defined belongsTo
    domain item
  slot-def manufacturedBy
    domain item
  slot-def printingSpeed
    domain printer
  slot-def printingTechnology
    domain printer
  class-def laserJetPrinter
    subclass-of printer
    slot-constraint
    printingTechnology
      has-value "Laser Jet"
  class-def HPPproduct
    subclass-of item
    slot-constraint
    manufacturedBy
      has-value HewlettPackard
  class-def HPLaserJet1100se
    subclass-of laserJetPrinter AND
    subclass-of HPPproduct
    slot-constraint printingSpeed
      has-value "8ppm"
```

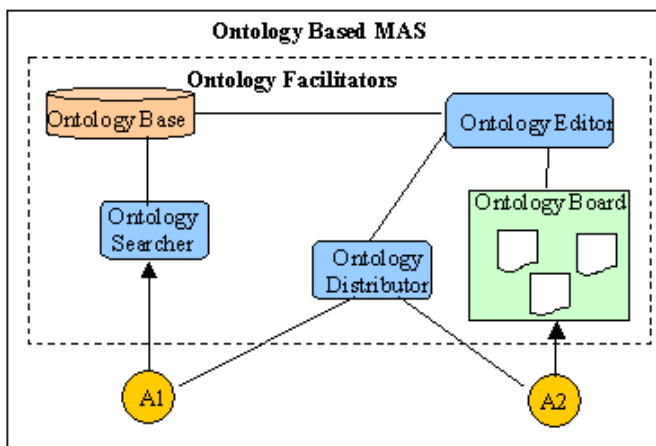
(2004b), we compared the advantages of some widely used ontology languages and illustrate how to use ontologies to describe knowledge in the domain of an online auction system. In Figure 1, we give a simple example to show how to describe an “item”, which is a kind of knowledge in online auction system domains, in an ontology. In this example, the ontology is expressed by the OIL language.

Ontology Facilitators in Ontology-Based MASs

For an ontology-based MAS, it is necessary to include ontology facilitators, which assist agents to search, acquire, and edit ontologies, into the system structure. We have introduced the necessary ontology facilitators, which should be included in an ontology-based MAS, together with the framework of ontology-based MASs in Bai and Zhang (2003; 2004a,b). In general, an ontology-based MAS should include the following ontology facilitators (Figure 2):

- **Ontology base** — the database that stores the common ontologies of the MAS;
- **Ontology board** — the mediator that receives and collects new ontologies from agents of the MAS;
- **Ontology editor** — the facilitator that checks and collects new ontologies from the ontology board, modifies and edits these new ontologies according to a standard ontology format of the MAS, then edits these new ontologies to common ontologies that can be read by all agents of the MAS;
- **Ontology searcher** — the ontology facilitator that searches related ontologies from the ontology base according to queries of agents of the MAS; and
- **Auction distributor** — the facilitator that publishes fresh or modified ontologies, which are received from the ontology editor, to related agents of the MAS.

Figure 2. *Ontology facilitators of ontology-based MAS's*



COLOURED PETRI NETS AND AGENT INTERACTIONS

Ontologies and ontology facilitators enable agents to include knowledge-level messages in their interactions. An inherent problem is how to enable agents to establish their interaction protocols flexibly according to their status. In this section, we present a CPN-based approach to flexibly form interaction protocols. We briefly introduce the basic principles of CPNs in the first sub-section, and explain how to apply CPNs to agent interactions in the second sub-section.


Petri Nets and Coloured PN's

The basic structure of a PN can be formally defined by the 4-tuple of Figure 3, where P is a set of *Places*, such as $P1$, $P2$, $P3$, and $P4$; T is a set of *Transitions*, such as $T1$, $T2$, and $T3$; A is a set of *Arcs*, such as the arc from $P1$ to $T1$, $T1$ to $P2$, $P2$ to $T1$; N is a set of *Tokens*. For example, in Figure 3, $P1$ and $P3$ have one token in the initial state. Net structure and transition firing rules are associated together to describe transfers from one system state to another. There are a number of transition firing rules associated with different types of PN's. However, all kinds of PN's share a common firing property: A transition can be fired if the token number of all input places is equal to or greater than their arcs' weights (Peterson, 1981). After a transition is fired, the tokens at its input places will be moved to its output places.

A CPN can be defined by a 9-tuple $(\Sigma, P, T, A, N, C, G, E, I)$ (Jensen, 1992), where Σ is a set of non-empty types, also called coloured sets; P is a set of places; T is a set of transitions; A is a set of Arcs; N is a node function; C is a colour function; G is a guard function; E is an arc expression function; and I is an initialization function.

CPNs differ from PN's because their tokens are not simply blank markers, but have data associated with them. A token's *colour* is a schema or specification. Places of CPNs contain *multi-sets* of tokens. Arcs of CPNs specify the token(s) that they can carry, and they can also specify some transfer conditions. Arcs exiting and entering a place can have an associated constraint function to determine which multi-set elements are to be removed or held. Transitions of CPNs are associated with guard functions that enforce some constraints on tokens.

Figure 3. Petri net example



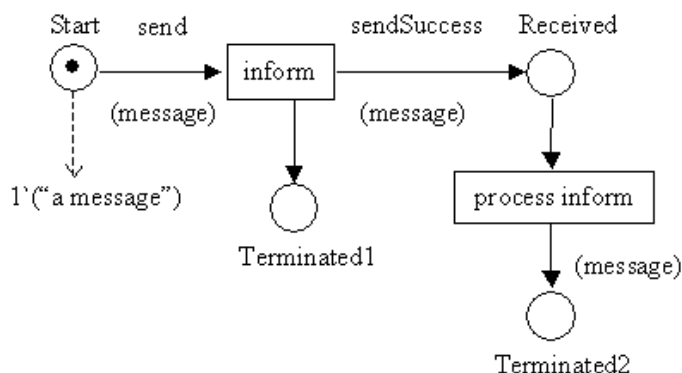
Using CPNs to Model Agent Interaction Protocols

There is a growing consensus that a CPN is a good way to model agent interaction protocols (Craneffield et al., 2002; Cost, 1999; Nowostawski, Purvis, & Craneffield, 2001; Poutakidis, Padgham, & Winikoff, 2002). By using CPNs, an agent interaction protocol can be modelled as a net of components, which are carriers of both protocol and interaction policy.

In using CPNs to model an agent interaction protocol, the states of an agent interaction are represented by CPN places. Each place has an associated type determining the kind of data that the place may contain. Data exchanged between agents are represented by tokens, and the colours of tokens indicate their data value. The interaction policies of a protocol are carried by CPN transitions and their associated arcs. A transition is enabled if all of its input places have tokens, and the colours of these tokens can satisfy constraints that are specified on the arcs. A transition can be *fired*, which means the actions of this transition can occur, when this transition is enabled. When a transition occurs, it consumes all the input tokens as computing parameters, conducts a conversation policy, and adds new tokens into all of its output places. After a transition occurs, the state (marking) of a protocol has been changed, and a protocol will be in a terminal state when there is no enabled or fired transition.

Here we take the FIPA inform protocol (FIPA, 2004) as an example of how to use CPNs to model agent interaction protocols. The FIPA request protocol can be modelled as a CPN, as shown in Figure 4. From Figure 3, we can see that there are five states in the interaction, and these states are represented in five places, respectively. If there is one or more tokens in the “Start” place and these tokens can satisfy the constraints specified in the “send” arc (belonging to data type “message”), the “inform” transition will be enabled. After the “inform” transition is fired, a token will be removed from the “Start” place, and the “Received” and the “Terminated1” place will get new tokens. After the “Receive” place gets a token, a “process inform” transition will be enabled. Finally, the interaction will be terminated after the transition “process inform” is fired.

Figure 4. Use CPN to represent the FIPA inform protocol



A CPN-BASED APPROACH FOR MULTI-AGENT INTERACTION

As mentioned in the first section, most agents are hard-coded within interaction protocols. This feature reduces the flexibility of agent interactions because protocols are hard to modify at runtime once they are pre-coded into the agents. In this section, we present a CPN-based approach to enable agents to flexibly form interactions. In implementing this approach, agents do not need to interact with other agents with a fixed protocol. Agent interaction protocols will be generated during interactions, and they can also modify their protocols according to their status during interactions. In general, the approach is composed of two main procedures, namely, sending protocol specifications and protocol analysis.

The Default Protocol and Sending Protocol Specifications

In this approach, agents of a system have the default interaction protocol shown in Figure 5. This default protocol comprises three components: two places that belong to the protocol specification (PS) type, and one transition. The related description of the default interaction protocol is described as follows:

- **PS (protocol specification)** — PS is a data type. A PS token contains a protocol specification that indicates an interaction protocol. The format of this specification is given in Figure 6.
- **CI (call interaction) place** — In the default protocol, a CI place is the place that conducts PS tokens. Agent may put a PS token, which specifies its desired interaction protocol, into a CI Place when the agent needs to interact with some other agent(s).
- **PS-set (protocol specification set) place** — In the default protocol, a PS-set place is the place that receives PS tokens, sent by other agents.

Figure 5. Default interaction protocol




Figure 6. Example specification of the inform protocol

```


/* InteractionProtocol
   Description: agent inform protocol */
(protocol
  (name inform)
  (place (name messagePlace) (owner sender) (PT message) )
  (place (name receivePlace) (owner receiver) (PT message) )
  (place (name failPlace) (owner sender) (PT error_message) )
  (transition (name sendMessage) (arc (name send)
    (from messagePlace) (match message) ) (arc (name
    sendSuccess) (to receivePlace) (match message) )
    (arc (name sendFail) (to failPlace) (match
    error_message) ))

```

- **SPS (send protocol specification) transition** — In the default protocol, a SPS transition sends protocol specifications (PS's). A CI place and PS-Set place are the input and output places of an SPS transition, respectively. The SPS transition can be enabled when there is/are token(s) in its CI place. A PS token will be removed from the CI place to the PS-Set place after the SPS transition fires.

The first procedure of the approach is the sending protocol specification. When an agent needs to interact with some other agent, it composes a PS to describe its desired interaction protocol according to its requirement, and puts the PS into its CI place. The content of PS is a specification of a CPN-modelled interaction representation. In this specification, the agent can indicate its required data as a place that needs the other agent to input a token(s). The data type constraints can be described in corresponding place types, and the data value constraints (token colours) can be described in constraint functions of arcs. The format of the specification can be as shown in Figure 6, or as defined by users. The corresponding CPN model is shown in Figure 7.

Figure 7. CPN model of inform protocol



Protocol Analysis

After an agent receives a protocol specification, it will analyse whether it can execute the interaction with its current status. In Petri Net theory (Peterson, 1981), there are many methods to analyse whether a PN model is executable. In this chapter, we use matrix equation methods (Peterson, 1981) to evaluate interaction protocols that an agent receives. Before introducing the protocol analysis method, we first provide some related definitions:

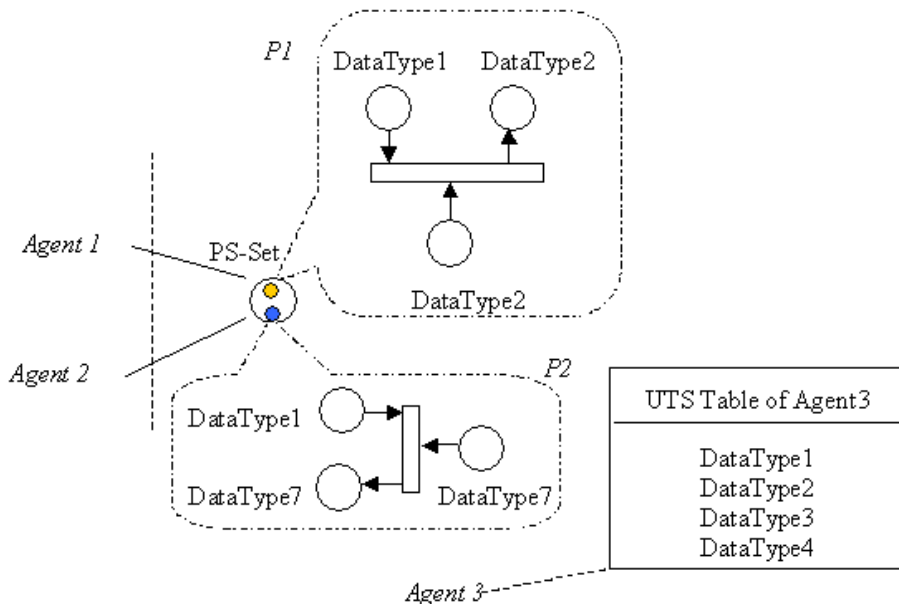
- **Definition 1:** The place type (PT) of a place is the associated data type that determines the kind of data/tokens which the place can contain.
- **Definition 2:** When an agent receives a PS, the PS place type set (PTS) is the set of place types that occur in the PS. For example, the PTS of the PS in Figure 7 is {message, error}.
- **Definition 3:** The understandable type set (UTS) of an agent is the set of data types that exist in the knowledge base of the agent.

Place Type Analysis

The first step of protocol analysis (place type analysis) is to test whether the agent can understand the PS that it received.

- **Definition 4:** If an agent with UTS receives a PS, the non-understandable type set (NTS) of the agent is: $NTS = PTS - (PTS \cap UTS)$, where PTS is the PS place type set of the PS and UTS is the understandable type set of the agent.

Figure 8. Place type analysis example



- **Definition 5:** An agent can understand a PS when $PTS \subseteq UTS$ or $NTS = \Phi$, where PTS is the PS place type set of the PS, UTS is the understandable type set of the agent and NTS is the non-understandable type set of the agent.

During place type analysis, the agent generates the NTS according to its UTS and the PTS of the PS token. If the agent can understand the PS, which means that the generated UTS is empty, the protocol analysis will process the second step. Otherwise, the agent will attach its comments, which indicate the place type that it does not understand, to the PS and send the PS back to the sender. In Figure 8, we give a simple example of place type analysis, where there are two PS tokens in the PS-set place (refer also to Figure 5) of Agent 3. These two PS tokens are P1 and P2, which are received from Agent 1 and Agent 2, respectively. The table lists the UTS of Agent 3. Comparing P1 and P2 with the UTS of Agent 3, we can see that Agent 3 can understand P1 but not P2 because Data7 in P2 is not in the UTS of Agent 3. Therefore, Agent 3 will comment, 'Data7 is un-understandable', on P2, and send P2 back to Agent 2.

Interaction Analysis

The second step of protocol analysis (state checking) is to test whether the current status of the agent is satisfactory to accept the PS, and whether the interaction will conflict with the goal of the agent. According to the matrix equation method (Peterson, 1981) of PN theory, a PN model can be expressed in a matrix format. For instance, the matrix format of the PN model of Figure 3 can be described by Equation (1), in which D^+ and D^- are matrices representing the output and input functions of the PN model, respectively.


$$D = D^+ - D^- = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} - \begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} 0 & -1 & -1 & 0 \\ 0 & +2 & +1 & -1 \\ 0 & 0 & -1 & +1 \end{bmatrix} \quad (1)$$

The marking of a PN model can be represented as marking set (m_1, m_2, \dots, m_n) , where m_n represents the token number of the n^{th} place. For example, the marking of Figure 3 can be represented as $(1, 0, 1, 0)$. In the context of interaction analysis problems, we introduce the following definitions, based on the matrix equation method:

- **Definition 6:** The *interaction matrix (IM)* is the matrix of an interaction protocol in PN model. IM^- and IM^+ are matrices used to represent input and output functions of the PN model. $IM = IM^+ - IM^-$.
- **Definition 7:** The *required token set (RTS)* of a PS is the *multi-set* of tokens that the PS requires the agent to offer.
- **Definition 8:** The *gain token set (GTS)* of a PS is the *multi-set* (Jensen, 1998) of tokens that the agent gets through interactions.

With the IM of an interaction protocol, an agent can calculate the GTS that it will gain through the interaction. Furthermore, the agent is able to check whether the perspective result of the interaction is in conflict with its own objective. For example, if Agent-1 received *Protocol-1* (as shown in Figure 9), the IM would be:

Figure 9. Gain analysis example



$$IM = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 2 \end{bmatrix} - \begin{bmatrix} 1 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 0 & 0 & +1 & -1 & 0 \\ 0 & 0 & -1 & 0 & 2 \end{bmatrix} \quad (2)$$


To make the transitions of Protocol-1 fire, we need to have a marking $m = (1, 1, 0, 1, 0)$. Since P1 and P2 have tokens in initial state, Agent-1 only needs to input tokens in P4. Therefore, we can have $RTS = (0, 0, 0, 1, 0)$. Furthermore, we can find out the GTS of Protocol-1 through the following calculations.

$$\mu' = \mu + x_1 \cdot IE = (1, 1, 0, 1, 0) + (1, 0) \cdot \begin{bmatrix} 0 & 0 & +1 & -1 & 0 \\ 0 & 0 & -1 & 0 & +2 \end{bmatrix} = (1, 1, 1, 0, 0) \quad (3)$$

$$\mu'' = \mu' + x_2 \cdot IE = (1, 1, 1, 0, 0) + (0, 1) \cdot \begin{bmatrix} 0 & 0 & +1 & -1 & 0 \\ 0 & 0 & -1 & 0 & +2 \end{bmatrix} = (1, 1, 0, 0, 2) \quad (4)$$

In the above calculations, Equation (3) shows the marking transition after T1 is fired, and Equation (4) shows the marking transition after T2 is fired (see also Figure 10). Since Agent-1 can only gain tokens from P5, the GTS of Protocol-1 is $(0, 0, 0, 0, 2)$. According

Figure 10. Gain analysis example (after T1 and T2 are fired)



to the data type of P4 and P5, we can see that through the interaction, Agent-1 will lose one token of DataType1 and gain two tokens of DataType2. Therefore, Agent-1 will evaluate whether the interaction is advantageous or harmful to its own goals.

POTENTIAL APPLICATIONS

In the previous three sections, we introduced how to use ontology to describe MAS knowledge and a CPN-based approach to form multi-agent interactions. These two methods are suitable for applications in complex dynamic environments. A potential application of the approach is supply chain formation (SCF) (Walsh, Wellman, & Ygge, 2000). A supply chain is a network that describes interrelated exchange relationships among multiple levels of production. SCF is the process of assembling complex production and exchange relationships between companies. To adapt to rapidly changing market conditions, companies need automated support for SCF to form and dissolve business interactions dynamically.

Agent technologies are widely applied in automotive supply chain formation. In such applications, the domain knowledge is usually mass and dynamic. Depending on the market conditions, factors such as produce varieties, price, and supply-demand relations are changeable. In this case, using ontologies to describe domain knowledge and including ontology facilitators in the MAS (refer to the section on *MAS Ontology & Knowledge Level Agent Interactions*) can bring lots of conveniences for knowledge acquiring.

Another challenge of SCF applications is how to coordinate finite resources and received interaction requests of agents. In SCF applications, agents may receive various interaction requests from other agents. On the other hand, some resources of agents are finite. A firm might be penalized if it accepts infeasible interaction requests. Therefore, the agent of a firm has to analyse received interaction requests and gives proper responses according to current resource availability within the firm. Using the CPN-based approach introduced in the third and fourth sections makes it easier for agents to analyse and dispose various interaction requests. For SCF applications, various kinds of resources and products can be defined as token data types, and a CPN can be used to describe supply-request (SR) relations of a firm. For example, the CPN of Figure 11 shows the SR relations of a firm called Firm-1. In this Figure, places P1 and P2 represent two kinds of products of the firm, R1 and R2 are received requests of P1 and P2, and S1, S2, ..., S5 are required resources to produce P1 and P2. At the current stage, Firm-1 accepts a request to produce P1, and S1, S2, and S3 all contain tokens. Therefore, the request can be satisfied. If Firm-1 receives another interaction request at this moment, it will analyse the received interaction protocol and make a decision according to its current status. Supposing Firm-1 receives the three different protocols in Figure 12, where Protocol-A requests the firm to supply product P2 and promises to requite M1; Protocol-B requests P2, promises to supply S4 and requite M1; and Protocol-C requests P2, promises to supply S3 and S4 and requite M1. According to the current status of Firm-1, Protocol-A is infeasible because of the shortage of resource S4; Protocol-B is also infeasible because resource S3 is occupied by request R1; only Protocol-C is feasible because the requester promises to supply S4 and S3. The above example can also be analysed by using the protocol analysis method introduced in the fourth section.

Figure 11. Resource-production relation example





Figure 12. Product request example



CONCLUSION

The social ability of an agent is exercised in a multi-agent system. For MASs, predefined agent interaction protocols reduce the flexibility of agent interaction, especially in open environments. In this chapter, we have proposed an approach to enable agents to form knowledge-level interaction protocols flexibly. Furthermore, in this approach, agents can also analyse whether the received protocol is understandable, whether the interaction can be accepted with the current status of the agent, and whether the interaction conflicts with the agents' objectives. These features make agents able to select or generate suitable protocols to interact with each other under open working environments.

ACKNOWLEDGMENT

This research was supported by an International Linkage Grant from the Australian Research Council (contract number LX0346249).

REFERENCES

- Bai, Q., & Zhang, M. (2003, December 17-19). An ontology-based multi-agent auction system. *Proceedings of the International Conference on Intelligent Technologies*, Chiang Mai, Thailand (pp. 197-202).
- Bai, Q., & Zhang, M. (2004a, February 16-18). Agent coordination through ontology management. *Proceedings of the International Conference on Artificial Intelligence and Applications*, Innsbruck, Austria (pp. 141-146).
- Bai, Q., & Zhang, M. (2004b, July 12-14). SWOAM — A semantic Web-based online auction multi-agent system. *Proceedings of the International Conference on Intelligent Agents, Web Technologies, and Internet Commerce*, Gold Coast, Australia (pp. 329-336).
- Cost, R. (1999, May 1). Modelling agent conversations with coloured petri nets. *Working Notes of the Workshop on Specifying and Implementing Conversation Policies*, Seattle, WA (pp. 59-66).
- Cranefield, S., Purvis, M., Nowostawski, M., & Hwang, P. (2002, July 15). Ontology for interaction protocols. *Proceedings of the 2nd International Workshop on Ontologies in Agent Systems — AAMAS'02*, Bologna, Italy. Retrieved July 18, 2004, from <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-66/oas02-16.pdf>
- Fensel, D., Harmenlen, F., Horrocks, I., McGuinness, D., & Patel-Schneider, P. (2001). OIL: An ontology infrastructure for the semantic Web. *IEEE Intelligent Systems*, 16(2), 38-45.
- Finin, T., Labrou, Y., & Mayfield, J. (1997). KQML as an agent communication language. In J. M. Bradshaw (Ed.), *Software agents* (pp. 291-316). Cambridge, MA: MIT Press.
- Foundation for Intelligent Physical Agents. *FIPA ACL Message Representation in String Specification*. Retrieved July 18, 2004, from <http://www.fipa.org/specs/fipa00070/>
- Horrocks, I., & van Harmelen, F. (2002). *Reference description of Daml+OIL ontology markup language* (Draft Report). Retrieved January 2002, from www.daml.org/2001/12/reference.html
- Jensen, K. (1992). *Coloured petri nets — Basic concepts, analysis methods, and practical use. Volume 1: Basic concepts*. Berlin, Germany: Springer-Verlag.
- Jensen, K. (1998). An introduction to the practical use of colored petri nets. *Lectures on Petri Nets II: Applications. Lecture Notes in Computer Science*, 1492 (pp. 237-292). Berlin, Germany: Springer-Verlag.
- Lesser, V. (1998). Reflections on the nature of multi-agent coordination and its implications for an agent architecture. *Autonomous Agents and Multi-agent Systems*, 1.1, 89-111.
- Lesser, V. (1999). Cooperative multi-agent systems: A personal view of the state of the art. *IEEE Transactions on Knowledge and Data Engineering*, 11(1), 133-142.
- Nowostawski, M., Purvis, M., & Cranefield, S. (2001, May 28-June 1). A layered approach for modelling agent conversations. *Proceedings of the 2nd International Workshop on Infrastructure for Agents, MAS, and Scalable MAS*, Montreal, Canada (pp. 163-170).

- Peterson, J. (1981). *Petri net theory and the modeling of systems*. Englewood Cliffs, NJ: Prentice-Hall.
- Poutakidis, D., Padgham, L., & Winikoff, M. (2002, July 15-19). Debugging multi-agent systems using design artefacts: The case of interaction protocols. *Proceedings of the 1st International Joint Conference on Autonomous Agents and Multi Agent Systems*, Bologna, Italy (pp. 960-967).
- Rao, A. S., & Georgeff, M. P. (1992, October 25-29). An abstract architecture for rational agents. *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, San Mateo, CA (pp. 439-449).
- Stanford KSL Ontolingua Server. (n.d.). Retrieved from <http://www-ksl-svc.stanford.edu:5915/>
- Walsh, W., Wellman, M., & Ygge, F. (2000, October 17-20). Combinatorial auctions for supply chain formation. *Proceedings of the 2nd ACM Conference on Electronic Commerce*, Minneapolis, MN (pp. 260-269).

Chapter IV

Literacy by Way of Automatic Speech Recognition

Russell Gluck, University of Wollongong, Australia

John Fulcher, University of Wollongong, Australia

ABSTRACT

The chapter commences with an overview of automatic speech recognition (ASR), which covers not only the de facto standard approach of hidden Markov models (HMMs), but also the tried-and-proven techniques of dynamic time warping and artificial neural networks (ANNs). The coverage then switches to Gluck's (2004) draw-talk-write (DTW) process, developed over the past two decades to assist non-text literate people become gradually literate over time through telling and/or drawing their own stories. DTW has proved especially effective with "illiterate" people from strong oral, storytelling traditions. The chapter concludes by relating attempts to date in automating the DTW process using ANN-based pattern recognition techniques on an Apple Macintosh G4™ platform.

INTRODUCTION: SPEECH PRODUCTION

Generally speaking, the aims of automatic speech recognition (ASR) are twofold: firstly, to extract the salient features of the incoming speech signals, then secondly to map these into the most likely word sequences, with the assistance of embedded acoustic and language models (Huang, Acero, & Hon, 2001).

Natural, conversational, continuous speech often incorporates false starts, repeated phrases, non-grammatical phrases (*ums* and *ahs*), and pauses, which bear little relation to written (text) punctuation. Some characteristics of speech which make recognition, whether by humans or machine, difficult include: background noise levels, variations in speaker loudness, pitch, emphasis (stress), and speech rate, not only *between* different speakers (either from within the same culture or due to different dialects), but also on different occasions with the *same* speaker (for example, with or without a head cold). Even worse, we tend to make assumptions as to what words (phonemes) we *expect* to hear next, based not only on the context of surrounding words (phonemes), but also on cultural mores. Further, since there is not always a strong correlation between the acoustic properties of speech waveforms and the linguistic units that they represent, this can lead to ambiguous interpretation. Ambiguities can also arise due to the fact that similar-sounding words can have quite different meanings (homonyms); conversely, different-sounding words can have similar meanings (synonyms).

A person's fundamental frequency (number of vibrations per second) is a function of their vocal cord mass, and typically ranges between 50 and 250Hz for males, and roughly twice this frequency for females.

We generate speech (phones) using a combination of voice box, or larynx (the vibration source), lungs (energy or power source), vocal tract and nasal passage (resonant cavities), together with the articulatory organs (lips, teeth, tongue, jaws, cheeks, and alveolar ridge — that region in the roof of the mouth which makes contact with the tip of the tongue) (Masaki, 2000). The lips, teeth, tongue, jaw, and cheeks are all capable of changing the shape of the basic resonant cavity, thereby producing different sounds. For example, the lips are involved in the production of English vowels and the consonants /b/ and /p/; the teeth (and lips) in /f/ and /v/; the alveolar ridge in /d/, /n/ and /t/, and the cheeks in /b/ and /p/. Likewise, various constrictions in our air passageways produce different sounds (for example, /p/, /b/ and /f/). Furthermore, sounds can be produced either with the vocal cords vibrating, referred to as “phonation” or voiced (for instance, /g/, /m/, /z/), or without, in other words “voiceless” (such as /f, /k/, /p/, /s/, /t/) (Keller, 1994).

Thus from a signal processing point of view, we can regard speech as a time-varying sound wave, whose frequency components are determined by changes in the size and shape of the vocal tract and associated physiology. Peaks in the energy spectrum of the speech waveform are referred to as acoustic resonant frequencies or “formants”. Most vowels comprise more than three formants; however, the first three (F1 ~500Hz, F2 ~1800Hz, F3 ~2500Hz), usually suffice for purposes of classification and/or recognition (higher-frequency formants reflect voice quality and individual speaker characteristics) (Ainsworth, 1997). Thus we can conceive of speech as the superposition of a number of frequency components of varying amplitudes and phases. As such, and in common with signal processing generally, speech is amenable to either Fourier series analysis (for continuous — analog — signals), or once digitized, to Fourier transforms (for discrete

signals). Speech recognition is invariably implemented on some form of computer platform; thus the raw speech signal must first be converted from analog to digital form.

Acoustic signals, including speech, are characterized by features such as pitch, duration, amplitude (loudness, signal strength, power/energy), and phase of each frequency component. As it happens, only the first three are relevant from a speech recognition perspective, since the human ear is insensitive to phase (Katigiri, 2000). Now since phonemes, the basic linguistic unit, are characterized by frequency, time, and energy, it makes more sense to use three-dimensional spectrograms rather than process the raw (albeit filtered) time-varying speech waveform. Filtering is necessary since speech, like any other one-dimensional time-varying acoustic signal, is susceptible to interference from background noise.

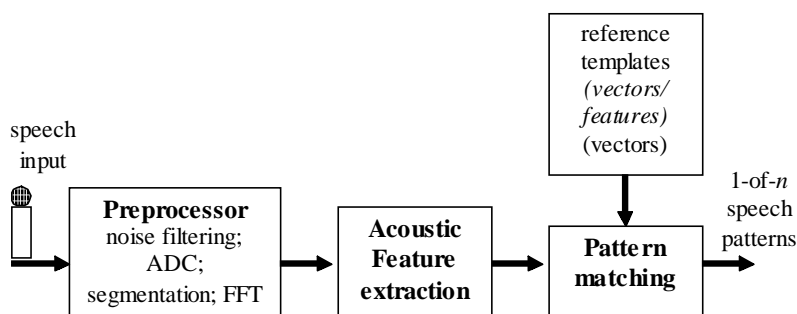
SPEECH RECOGNITION

Humans use not just auditory information in recognizing speech, but a host of non-verbal cues as well — more specifically, a speaker's facial movements (mouth, eyebrows, and so on), body gestures, the direction from which the sound arrives, background noise levels, cultural context, knowledge of the language (vocabulary), dialect/accent, the influence of the local environment (and more globally the state of the world at large), and so forth. Prosody refers to longer timescale (supra-segmental) speech characteristics which assist listeners to distinguish between utterance types (namely, statements/declarations *versus* questions/interrogations *versus* commands/imperative sentences, in other words, the speaker's *intent*), to clarify both sentence structure and syntax, to interpret word emphasis (stress, timing, rhythm, melody, and/or intonation), to indicate whose turn it is to speak in the conversation (pauses), and to correlate acoustic structure with the speaker's emotional state (Cosatto, Ostermann, Graf, & Schroeter, 2003; Keller, 1994).

We also need to be cognizant of the fact that the ear canal itself constitutes a resonant cavity, and as such, filters sounds prior to their impacting on the ear drum, and thence, via the auditory nerve, to the aural processing area in the brain.

All of this makes the automatic recognition of continuous speech by computer a difficult task. Further, a lot of important human speech processing takes place between the ear drum and the auditory nerve, prior to reaching the aural processing region of the brain.

Figure 1. Generic ASR system



Isolated, single-speaker recognition is a much more viable proposition than continuous, multi-speaker recognition; the latter is a much less constrained problem, and as such is considerably more difficult to solve. For one thing, detection of inter-word gaps is much more of a challenge. More constrained tasks, such as isolated word recognition, word spotting or speaker authentication, more readily lend themselves to implementation by computer; indeed, numerous such applications have made their way to the marketplace (see the *State-of-the Art* section).

From a machine learning perspective, speech signals are regarded as yet another time-varying pattern. As such, standard pattern recognition techniques can be applied. Not surprisingly though, more specialized and appropriate methods have been developed for speech recognition.

A typical speech recognition system is shown in Figure 1.

Processing of speech signals can be undertaken either in the time or frequency domains. High pass filtering of a raw speech signal will produce RMS amplitude envelopes; application of a fast Fourier transform will facilitate the extraction of frequency components.

Automatic speech recognition has traditionally been attempted using template matching and distance measures, probabilistic classification, and/or artificial neural networks.

Pre-Processing

Prior to recognition proper, the speech signal must first be pre-processed, in order to filter out background noise, to correct for microphone distortion, and to convert from analog to digital form, in order to facilitate subsequent processing on a computer. The speech waveform is next translated from the time domain to the frequency domain by way of fast Fourier or similar transform.

Figure 2. 3-Dimensional spectrogram

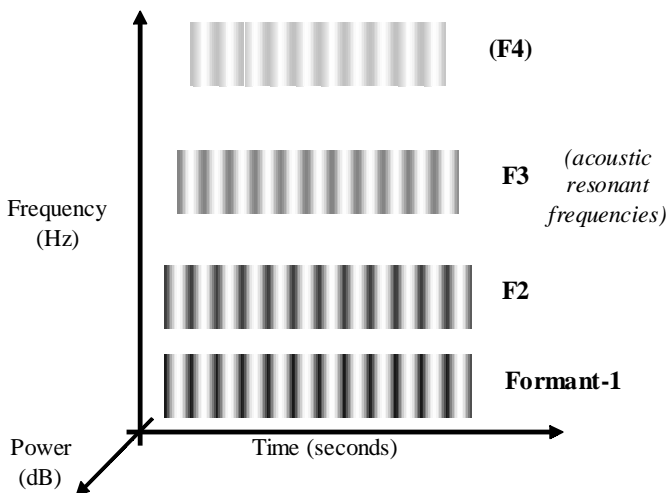



Figure 3. Short-timescale spectrum



The resulting 3D spectrogram (frequency *versus* time *versus* signal strength/intensity, see Figure 2) assists in the recognition process. Narrow-band spectrograms reflect changes in pitch (frequency), whereas wideband spectrograms provide information regarding formant structure. Furthermore, a 2D “slice” through this 3D spectrogram at a specific time yields a “short-time” spectrum (energy vs. frequency) characterized by constituent harmonics and an overall spectral envelope (Figure 3). Peaks in the former correspond to the various acoustic resonant frequencies (formants), whereas the latter reflects the shape/size of the various articulatory organs involved in producing the speech (namely, the lips, tongue, and so on). Formants are the principal determinants of vowels, while the harmonic structure is a result of vocal cord vibration, which is considered to mainly convey speaker identity and prosodic information (Katigiri, 2000).

Segmentation Methods

One of the primary pre-processing tasks is to separate the incoming (acoustic) speech signal into basic “units” or segments. This can be performed on the basis of time, various signal characteristics (such as fundamental frequency, overall spectral shape or signal strength (power)/energy threshold), or on the basis of phonetics (the science of speech). Moreover, the segment size can be either fine- or coarse-grained. For small vocabulary, meaning isolated word recognition (word spotting), large size segments suffice; for large vocabulary, meaning multi-speaker, continuous speech recognition, shorter timescale units are more appropriate.

It should be pointed out that in word spotting (or key phrase detection), the entire utterance is comprised of *known* words, and as such constitutes a much more constrained problem than with ASR generally (in which we do not know which words are coming next in sequence). In other words, the former is amenable to classic pattern recognition techniques.

Phonemes are the smallest discrete sound units (abstractions), allophones are variants of phonemes used in everyday speech which pertain to specific contexts, and diphones (dyads or demi-syllables) are transitions between neighbouring segments. By contrast, a syllable is a larger, separately utterable segment, and comprises a linear sequence of phonemes of the form consonant-vowel-consonant (CVC). For the purposes

of speech recognition, both vowels and consonants can be defined in terms of their phonetic properties; however, incorporation of prosody characteristics invariably improves performance. To put this into perspective, the English language comprises around 300,000 words, but only ~10,000 syllables, ~12,000 morphemes (grammatical categories), between 1,000 and 2,000 demi-syllables/diphones, and only around 40 phonemes (16 vowels + 24 consonants). Furthermore, *spoken* phoneme sequences bear little relation to the spelling and punctuation which characterizes *written* text (especially for English).

Acoustic Feature Detection and Extraction

Various acoustic features can be extracted from speech signals, whether these be prosodic, or alternatively derived from time- or frequency-domain characteristics, as previously observed. A number of these can be combined to form an acoustic feature vector, which in turn can be used for pattern matching/recognition/classification purposes, in other words, as reference templates (vectors) stored in a look-up table (Figure 4). Thus words can be regarded as a sequence of feature vectors. In some cases, these features are sufficiently distinct (well-defined) that they can be readily separated into classes that correspond to non-overlapping regions in the feature space. Hence the acoustic features can be grouped into categories (classes) without the need for training (supervision).

This raises the issue of the overhead needed to store such a look-up table. In this context, vector quantization, the process of categorizing input data into clusters (codebook vectors) prior to transmission/storage, can be useful. Vector quantization divides n -dimensional space into smaller 2^m regions, where m is the number of codebook

Figure 4. n -Dimensional/element acoustic feature vector






Figure 5. Acoustic filter bank



address bits, known to both transmitter and receiver. This codebook can be used to translate vectors in n-D space (actually the centres-of-gravity of “regions”) into addresses with minimum distortion, and vice versa. Thus, VQ can be thought of as performing data compression (dimensionality reduction) prior to storage. Alternatively, VQ can be regarded as performing K-means clustering on the incoming speech signals (Sridharan, Leis, & Paliwal, 2000).

Extraction of speech features can be undertaken using a variety of different methods, including acoustic filter banks (Zwicker, 1961), auto-regressive modeling (Atal & Hanauer, 1971; Itakura & Saito, 1970), cepstrum, the inverse Fourier Transform of the

Figure 6. Linear predictive coding (auto-regressive modelling)



logarithmic spectrum (the peak of which corresponds to pitch), dynamic feature modeling (McDermott & Katagiri, 1991; Waibel, Hanazawa, Hinton, Shikano, & Lang, 1989), and probabilistic methods. In the early days of ASR, filter banks (Figure 5) were popular. These gave way during the 1970s and early-to-mid 1980s to linear prediction coefficients (LPCs) (Figure 6). Since that time, the cepstrum approach has been predominant, since it leads to less variability. For instance, in the HTK System, 12 mel-frequency cepstral coefficients (MFCC), together with the signal strength, are combined to form a 13-element acoustic feature vector (Young, Wood, & Byrne, 1996).

Linear predictive coding is a parametric method, in which the next sample in the time series is derived from a combination of the weighted sum of the previous (and therefore known or predictable) samples X_{t-i} , together with unpredictable noise, as indicated in Eqn. 1:

$$X_t = -\sum_{i=1}^n \alpha_i X_{t-i} + e_t \quad (1)$$

The coefficients α_i can be estimated using either auto-correlation or covariance techniques. The set of LPCs for each sampling time interval constitutes an acoustic feature vector.

Dynamic Time Warping

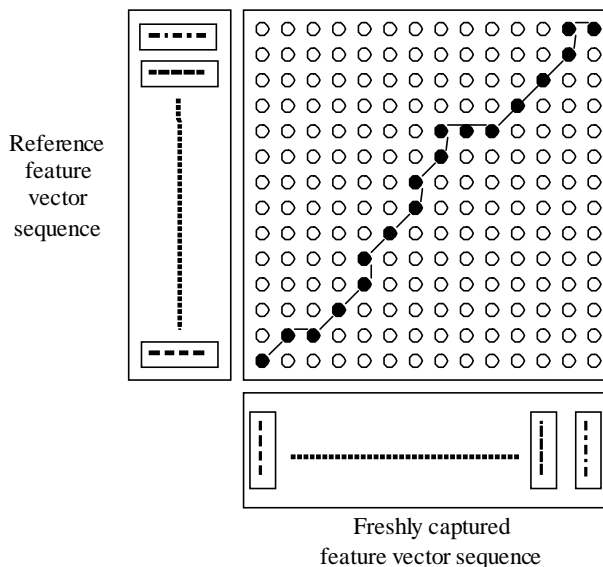
Since we regard speech signals as just another type of pattern, one approach to ASR is to compare new utterances with templates stored in a look-up table. This is not as straightforward as it first appears, however, since people speak at different rates (including the *same* speaker on *different* occasions). This means that such stored templates will need to be “stretched” or “compressed” prior to matching. This is precisely what is behind the idea of dynamic time warping.

Short (10-30mSec) segments are typically used in dynamic time warping. In Figure 7, a freshly-captured sequence of speech vectors is compared with a reference vector sequence (template) stored in a look-up table. Obviously this will only be effective for single speakers; multi-speaker recognition would necessitate the use of *averaged* reference feature vectors, and hence would lead to degraded performance.

A dynamic programming lattice is maintained during successive iterations (sample periods), indicating how close (in a minimum distance, dissimilarity, or short-time spectral distortion sense) each feature vector is in the sequence; obviously a 45° straight line indicates a *perfect* match. Thus, dynamic time warping automatically adjusts the time scale of the incoming speech signal to better align with the stored templates prior to matching proper.

Dynamic time warping is effective for isolated word recognition, but becomes less viable as the number of templates increases, due to the exponentially-increasing time required to compute the distance matrices. Notwithstanding, the dynamic time warping technique *can* be extended to handle sequences of words, and even continuous recognition, but with the latter we also need to consider syntax in order to discount illegal word sequences.

Figure 7. Dynamic time warping example



Hidden Markov Models (HMMs)

Dynamic time warping uses template matching; by contrast, HMM uses a probabilistic approach based on Bayes decision theory, and generally speaking, leads to better speech recognition performance. Any stochastic (random) process can be represented by an n th-order Markov model. Hidden Markov models can be thought of as stochastic systems (automata) which comprise a state vector- \mathbf{S} , a transition matrix- \mathbf{V} , and an emission matrix, the latter comprising the probabilities of “emitting” a particular “symbol” from a given state- S_i . In other words, each state has an associated probability distribution function (PDF) $P(\mathbf{V}|S_i)$ for each acoustic (feature) vector \mathbf{V} , together with a set of state transition probabilities $P(S_i|S_j)$, in other words, of moving from state- S_i to state- S_j (Ferguson, 1980; Huang, Akiri, & Jack, 1990). A simple two-state HMM is shown in Figure 8. During training, these probabilities are estimated from the sequences of symbols produced by typical (representative) speech utterances, using the Baum-Welch algorithm, which correlates probabilities with local maxima.

Recognition requires much less computation than training. It is possible to use a recursive formula to calculate the probability that each word HMM produced the observed data, the model with the highest probability of being selected as the one to best represent the word in question. An efficient dynamic programming technique, the Viterbi algorithm, can be invoked to efficiently compute the probabilities of *each* HMM generating a specific spoken sequence of speech symbols, in other words, the most likely path through the various model states (actually, these probabilities will invariably be less than the *true* probabilities; nevertheless, in most cases it leads to acceptable recognition performance).

Figure 8. A two-state hidden Markov model


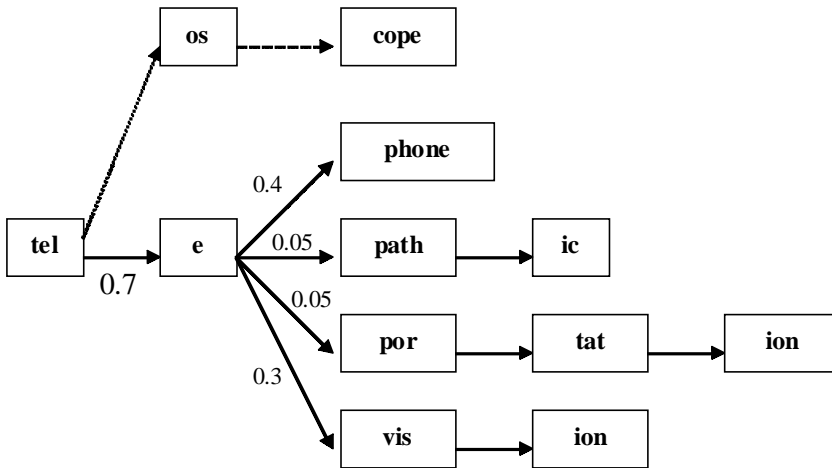


Figure 9. HMM (written syllable) example



Since some word segments (phonemes) exhibit more variability than others, it is possible to generate a stochastic model for *every* word in our reference vocabulary (for example, Figure 9). By contrast, an example of an *invariant* rule (albeit with written rather than spoken English) is that a “q” is invariably followed by a “u”; by contrast, *many* different letters can follow the letter “b”.

In the context of speech recognition, HMMs can be used to represent different words as a different sequence of states, together with probabilities of moving from one state to permitted successor state, along with probability distributions defining the expected observed features for each state. By the way, the “hidden” in HMM refers to the fact that only the acoustic vectors are observed, not the state sequences.

HMMs lead to improved performance compared with dynamic time warping, and indeed is the *de facto* speech recognition method in use nowadays.

Hybrid ANN/HMM systems estimate the class condition probabilities $P(\mathbf{V}|S_i)$ of a speech vector \mathbf{V} , given that the system is in state S_i , using ANNs rather than Bayesians

(Katagiri, 2000). Since ANNs learn *a posteriori* probabilities — $P(S_i|V)$ — they must first be converted to class conditional probabilities before they can be used in HMMs. This can be achieved using Bayes' theorem:

$$P(V|S) = \frac{P(S|V)P(V)}{P(S)} \quad (2)$$

Artificial Neural Networks (ANNs)

Despite their well known limitations, artificial neural networks (ANNs) are admirably suited to pattern classification and/or pattern recognition (Fulcher, 1997). They would appear, then, to be a natural approach for speech recognition, since from a classification/recognition perspective speech is simply another pattern (Lippmann, 1989). As it happens, and despite a lot of research effort during the 1980's and 1990's, the *de facto* speech recognition technique used nowadays is not ANNs, but HMMs. Nevertheless, as we shall see, ANNs are especially suited to the application of interest in this chapter, namely, speech recognition for developing literacy in non-text literate users.

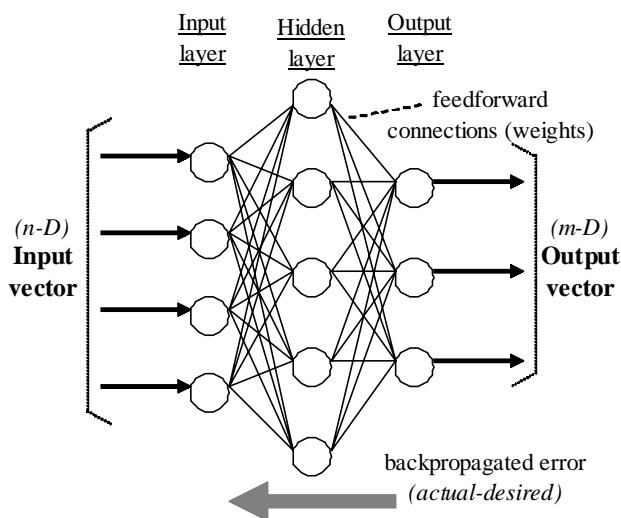
Several different types of ANN have been applied to speech recognition over the years, including multi-layer perceptron/back-propagation (Brunet, Pandya, & Pinera 1994; Wu & Chan, 1993), Self-Organizing Map (Kohonen, 1997), Time-Delay Neural Networks (Lang, Waibel, & Hinton, 1990; Waibel, Hanazawa, Hinton, Shikano, & Lang, 1989), recurrent networks (Robinson, Hochberg, & Renals, 1996; Schuster, 2000), Learning Vector Quantization (McDermott, 2000; McDermott & Katigiri, 1991), and probabilistic neural networks (Haton, 2000), not to mention modular/hierarchical ANNs (Jordan & Jacobs, 1992, 1994) and hybrid systems (Fritsch, Hild, Meier, & Waibel, 2000). We focus here on MLP, SOM, and TDNN.

MULTI-LAYER PERCEPTRON/ BACK-PROPAGATION (MLP/BP)

When a lay person refers to ANNs, they are invariably referring to multi-layer perceptrons trained using the back-propagation learning algorithm. A typical such network is shown in Figure 10.

Representative input-output training (vector) pairs are presented to the network in succession. For each presentation, the error difference between the *actual* output and the *desired* output is used to alter the network weights, firstly those connecting the hidden layer to the output layer, thence those connecting the input and hidden layers. In this manner, the errors are said to propagate *backwards* from output layer to input layer, hence the term Back-Propagation learning algorithm. Repeated presentations of *all* input-output training exemplars is needed, since the weights will be adjusted in many different directions by the time the first exemplar pair is revisited. Not surprisingly then, many iterations are normally required in order for the network to converge to a global minimum, such that *all* input-output pattern pairs have been learnt. Consequently, one inherent feature of MLPs is their long training times; nevertheless, once trained, such

Figure 10. Multi-layer perceptron/back-propagation



networks can respond instantaneously to inputs they have not previously met (and recall the output pattern associated with it). As is common with all ANNs, MLPs are capable of generalization, and moreover are both noise- and fault-tolerant.

Based on the Kolmogorov representation theorem, it can be readily shown that two hidden layers at most are required to provide arbitrary decision boundaries in the solution space (energy landscape). Furthermore, feed-forward networks with non-polynomial activation functions are capable of approximating any continuous function, to any degree of accuracy (Hornik, 1991; Leshno, 1993). Zhang, Xu, and Fulcher (2002) extended this important result to higher-order ANN groups, in which each element is a standard MLP/BP and uses piecewise rather than polynomial activation functions and thresholds; more specifically, such ANN groups can approximate any kind of piecewise continuous function, to any degree of accuracy.

Now from a speech recognition perspective, the speech signals are treated as just another input pattern, regardless of whether this vector is derived from digitized time-varying waveforms or from 3D Spectrograms (the frequency domain).

SELF-ORGANIZING FEATURE MAP (SOM)

Unlike MLPs, Kohonen's self-organizing feature map (SOM) is an *unsupervised* network. It is not trained using representative exemplars; rather, it forms its own classifications, which may or may not make sense to the user.¹

SOMs are a kind of associative memory inspired not only by the self-organizing and adaptive features of the human brain, but more especially by its localized activity. More specifically, SOM is based on the premise that the brain uses spatial mapping in order to internally map complex data structures.

Categories are formed within SOM, according to the statistical properties of the input patterns, as opposed to those imposed by the trainer in a *supervised* network. The input speech signal is fed into a 2D array, each node of which connects to its nearest neighbours, and which has a weight vector associated with it. The node which most closely matches the input (in a Euclidean distance sense) is selected, and its weight vector, together with the weight vectors of all nodes in its “neighbourhood”, is moved to more closely align with the input vector. During training, the size of this neighbourhood is gradually reduced, with its boundaries determined by the so-called Mexican hat function. These neighbourhoods act as feature classifiers on the input data. The nodes in the output layer start off as randomly organized (vectors pointing in random directions), but end up as a self-organized feature map.


The above competitive learning, or clustering, constitutes an ordered mapping, and an output is produced in the form of a 2D topological phoneme (phonic) map which corresponds to the input speech sequence; the phonic map produces a dynamic trace of the spoken words. What is more, continuous speech input patterns will trigger the *same* path through this 2D phonic map, independent of speech rate.

Kohonen (1997) used SOMs as the basis for his “phonetic typewriter”, which produces written text when presented with speech input. Actually, the phonetic typewriter uses a hybrid system, a combination of SOM and rule-based expert system, the latter being needed in order to resolve context due to the surrounding phonemes (grammar).

Output node clusters are formed during training and labelled manually afterwards, with the assistance of the expert system rule base.

The usual Nyquist Sampling criterion applies to the speech input, namely, that the sample rate needs to be at least twice the highest frequency component of the incoming speech signal. Pre-processing takes the form of 5.3 KHz low pass filtering, 12-bit analog-to-digital conversion, and 256-point FFT conversion. The resulting 15 frequency band

Figure 11. Self-organizing feature map



continuous (normalised) pattern vector, together with the RMS value of the speech signal, are combined to form the 16-bit feature vector fed into the SOM. Post-processing (as described earlier) is invoked to correct misclassifications.

The inputs to the SOM are time slices of the speech waveform; outputs are phoneme classifications, decided on the basis of several consecutive inputs.

Reported accuracy is between 85-95% for the Finnish language, which, unlike English, is largely phoneme-based.

RECURRENT AND TIME DELAY NEURAL NETWORKS

The problem of different speaker rates was addressed previously by way of dynamic time warping, which effectively compresses (“stretches”) the raw speech waveform in order to better match the timeframe of the stored (reference) template, prior to pattern matching proper. Alternative approaches include recurrent ANNs and time delay neural networks.

MLPs are only capable of learning *static* (time-independent) input-output (vector, pattern) mappings. In order to model *dynamic* systems, a neural network needs to incorporate some form of memory, in other words, prior knowledge. One way of achieving this is to add time delays (the values of which are modified during learning) to the basic MLP architecture, resulting in the recurrent network of Figure 12.

A sequence of acoustic input vectors presented to the network produces a corresponding sequence of output and internal state vectors. After presentation of the last vector, the final output and state vectors are compared with the target (desired) output and arbitrary state vectors, respectively. Error signals are derived, working in reverse from the last to the first vector in the sequence. The weights are then altered by an amount equal to the average weight change resulting from the preceding back-propagation phases. In this manner, temporal variations are reflected in the internal

Figure 12. Recurrent neural network





Figure 13. Time delay neural network



network states. Moreover, the path traveled through the state space is time-independent, unlike TDNNs which, as we shall see, employ a *fixed* time window.

Time delay neural networks are feed-forward ANNs whose hidden and output neurons are replicated across time (Lang, Waibel, & Hinton, 1991). Different parts of a TDNN perform the same computations but on different time-shifted versions of the inputs. The input to each node is a combination of the current input and the n previous inputs (where $n=2$ for H1, 4 for H2, and 7 for the output layer, as shown in Figure 13).

Waibel et al. (1989) used the TDNN of Figure 13 to discriminate between the isolated syllables “bee”, “dee” and “gee”.

In a related study, Lang and Hinton (1988) applied a similar TDNN to “bee”, “dee”, “ee” and “vee” discrimination (which obviously necessitated the use of four rather than three output neurons). They achieved a recognition accuracy of 93% on test, not training data. In the latter TDNN, 192 (16×12) input neurons encoded the 2D spectrogram, the hidden layer consisted of 10 copies of eight neurons, and the output layer comprised six copies of four neurons. The total number of synaptic weights was 544. The various hidden neuron replicas applied the same set of synaptic weights to narrow (three time step) windows of the spectrogram. Likewise, the various output neuron replicas applied the same set of synaptic weights to narrow (five time step) windows in the “pseudo-spectrogram” computed by the hidden layer.

Multi-stage TDNNs extend the basic idea from single phoneme recognition to phoneme sequences (words), and incorporate non-linear time alignment of the incoming speech with stored models (not unlike that described earlier with dynamic time warping) (Fritsch et al., 2000).

OTHER SOFT COMPUTING APPROACHES

Soft computing refers to techniques that are not algorithmic-based, with hard and fast inbuilt rules; rather, they involve some form of learning (iterating over time towards a viable solution), and are, in the main, biologically inspired. Artificial neural networks, genetic/evolutionary algorithms, genetic programming, and fuzzy logic are commonly grouped together under this banner. These techniques can also be used to construct hybrid systems, thereby leading to improved performance (Stadermann & Righoll, 2003). Other examples are the GA refinements to both DTW and HMMs (Man, Tang, & Kwong, 1999), and ANN/HMM hybrid approaches (Ganapathiraju & Picone, 2000; Trentin & Gori, 2004).

THE STATE-OF-THE-ART

Automatic speech recognition has made great strides over the past several decades, and most notably during the last few years due to successive refinements of the basic algorithms. Automated directory assistance, voice dictation, help desks, and even airline reservations have become commonplace, the performance of which, however, remains variable.

Despite this, the field is still viewed as rather fragile and immature, with industry and the public reluctant to embrace the available technology. Two major challenges remain with ASR in the large, these being reduction of: (a) the substantial difference in performance between noise-free and noisy acoustic environments, and (b) recognition error rates in conversational/casual speech (Deng & Huang, 2004). Multi-modal and multi-sensor system designs are being adopted in order to adequately address these challenges. Another driving force in the advancement of ASR technology is the annual DARPA-sponsored speech-to-text evaluation program, which has inspired a roughly 10% reduction per annum since its inception in the late 1980s (www.nist.gov/speech/tests/spk/2000).

On the other hand, speaker *authentication* technology is maturing at a rapid rate, with several products now on offer in the marketplace (for example, Scansoft's SpeechSecure™, Courion's PasswordCourier®, Vocent Solutions' Confirmed Caller™, and similar voice authentication offerings from Nuance Communications, Voice.Trust, VoiceVault, and similar start-ups) (Vaughan-Nichols, 2004). It should be pointed out that speaker/voice authentication/verification is a much more constrained problem compared with ASR and/or natural language understanding, and hence more amenable to solution. Nevertheless, and in common with *all* biometric techniques, voice authentication faces considerable hurdles to widespread adoption by the general public, not the least of which are concerns over security and privacy (Fulcher, 2005).

Compared with alternative biometric approaches, voice authentication possesses certain advantages and disadvantages. The former include less intrusiveness and the ability to perform *remote* authentication; the latter include lower accuracy, greater template storage requirements, and users' unwillingness to repeat themselves when prompted by the system. From a speech recognition perspective, the stored template (user "voiceprint") comprises frequency information, phoneme pronunciation, word combinations, accents, and additional linguistic/idiomatic features (such as pauses, ums

& ahs, and the like). Voice authentication has found its way into credit card products, such as ComDot™ (<http://www.beepcard.com>).

Voice translators have also started to migrate from the research labs to the marketplace, especially for mobile computing, for example, the Arabic-English Speechalator (CMU, Cepstrum, Mobile Technologies & Multimodal Technologies), Phraselator® (VoxTec), and NEC. The task of such translators is entire *phrase* rather than individual word recognition, and as such is less ambitious (and hence more doable).

Also, considerable effort has been put into the continued development of “industry standard” speech platforms such as the open source Sphinx suite of tools (Walker, Lamere, Kwok, Raj, Singh, Gouvea, Wolf, & Woepel, 2004), the World Wide Web Consortium’s (WC3) VoiceXML, and associated Speech Recognition Grammar Specification (SRGS), Speech Application Language Tags (SALT) by Cisco, Intel, and Philips, and Microsoft’s Speech Server technology (targeted for Visual Studio.NET™).

Lastly, futuristic speech research continues at leading institutions, including IBM (“superhuman” speech), Microsoft, Intel, and Carnegie-Mellon University, to mention but a few.

THE DRAW-TALK-WRITE APPROACH TO LITERACY

Draw-talk-write — DTW — is a process that enables literacy-inefficient, visually-strong and orally-proficient people to become literate in a range of contexts and disciplines, on some occasions enabling them to proceed on to tertiary studies (and successfully so) (Gluck, Vialle, Lysaght, & Larkin, 1998). The process has evolved through working with people who have had stories to tell and who have also had extreme difficulty in putting the story of their ideas, thoughts, work, and experiences into text in a form that meets the requirements of their audience, discipline, and most of all, themselves.

Evolution of the DTW process began as one of the authors (Gluck) worked with indigenous Australians in squatter camps, jails, pre-vocational and vocational courses, undergraduate and postgraduate degrees to enable them to work toward meeting their literacy-specific needs. The process has far-reaching repercussions beyond this community sector (Gluck, Vialle, & Lysaght, 1999a,b). DTW continues to evolve through collaborative work with both indigenous and non-indigenous Australians to facilitate the drawing, telling, and writing the stories of their PhD (and other Doctoral) dissertations. Fundamental to the DTW approach is that students’ own rich personal stories serve as the mediating agent. Moreover, DTW by way of speech recognition has underpinnings in Vygotsky’s concept of “tool mediation” (Rieber & Carton, 1987; van der Veer & Valsinger, 1994).

Getting Started: Asking for Help, Safety to Learn, and Joining In

The majority of students entering the DTW process do not feel safe to write. Their experiences of writing have generally been ones of struggle, not feeling up to the task,

and knowing themselves as “inadequate” writers. They have also felt, and in many cases experienced, what they have written as not “good enough” for the audience they have been writing for. Most significantly, what they have written has not met their *own* expectations. Students’ fear of failure has at times been palpable. To say that writing and reading has not been a safe experience for the overwhelming majority of such students is an understatement. They have risked asking for help and joining the program, often as a means of last resort, to try an alternative way which involves drawing, talking, writing then reading, and finally redefining or retelling their story of themselves as writers.

The students’ preparedness to ask for help and try another way has provided an opportunity for them to utilise the DTW method to recreate their stories of themselves as “adequate”, “good”, or even “excellent” writers. It also enables them to experience and know writing as a rewarding activity.

The challenge for the learning facilitator is whether they have what it takes to listen to, hear, acknowledge, and respond to students’ needs in a way that enables the latter to access and utilise the help on offer, in other words, to develop a safe place in which the student and facilitator can join and begin to journey from where they are to where they want and need to go (Gluck & Draisma, 1997).

DTW has evolved in the context of creating learning spaces within which students feel safe to ask for and receive help, and learning facilitators are safe to offer help and to put this in a form that can be accessed and used by the students, even in the face of opposition from “conventional” teaching practitioners and educators.

Determining Students’ Needs and How to Meet Them

Possibilities for determining what the students’ needs are and how these can be met are revealed by the students as they feel safe enough to risk beginning to join with the learning facilitator working on their writing assignments. For example, as they talk about what they are trying to write, the students begin to reveal snippets of their learning and writing stories through what they say and do. Their talking, together with the facilitator’s listening, observation, and interaction with the students, are essential parts of forming successful learning relationships, because as they talk and position themselves within the workspace, they begin to reveal their fears of writing and their desires to successfully complete the written assignment.

As students talk, they frequently begin to risk sharing their stories of writing and reading, and the facilitator begins to hear and gain an understanding of their experiences of themselves as writers. While the students talk, the facilitator listens to their voices, and checks what they have heard against observations of the students’ physical posture and their positioning of themselves in relation to the facilitator and the chaotic state of the physical learning space (the facilitator’s office is a conglomeration of incongruencies, with no books on the shelves, an electronic whiteboard that acts as a theatre scrim upon which to project and record images, an office floor which serves as a stage, a pink pig with wings sitting on top of a computer, and a purple tutu with silver trim adorning the door; as a colleague has remarked: “One *responds* to incongruencies”. More specifically, incongruencies can be used in the learning context to generate possibilities for both the students’ and the facilitator’s learning).

Where there is congruency between the student’s voice, body posture, and their positioning of themselves within the room, the facilitator is then able to identify and

respond to the learning needs embedded in their stories. Listening and observation is from the perspective of feeling one's way into what the students are saying, hearing the essence of their learners' voices, and then responding to those voices in a voice that is empathic, and furthermore is in language and actions that the students can identify and resonate with. Responses frequently employ humour and street language that students would normally use away from the formality of a "right and proper" teaching situation.

Listening to Align

Listening to learners' stories enables alignment with the essence of their voices, as well as identification of their learning strengths, or "dominant intelligence" (Gardner, 1983). Gardner defines intelligence as "the ability to solve problems, or to create products, that are valued within one or more cultural settings", and proposes seven distinct intelligences, these being: visual/spatial, musical, verbal, logical/mathematical, interpersonal, intrapersonal, and bodily/kinaesthetic. By gaining an understanding of the student's dominant intelligences, we are able to reframe positions and actions in the learning context and process so that they are directed toward facilitating a student's individual and collective learning through their dominances and gifts. The focus is to utilise the student's strongest ability to bring their stories "in-line" with other modes through which clear communication with others is most easily achieved. In turn, this provides a medium and way of being and acting with students that brings them, their learning story, their dominances, the assignment in question, the relevant theory, and inherent literacy demands of the disciplines they are studying, and the facilitator into relationship.

Effectively understanding student strengths provides a means of facilitating "conditions for learning" that enable the student to be immersed in the discipline, and to receive demonstration of how learning and assessment are structured in the culture of the discipline (Cambourne & Turbill, 1987). This provides them with a basis for developing expectations of what "success" means in the discipline, and to come to grips with how they can begin to respond to what is required. Demonstration and expectations enable the students to do away with "guess work" and the idea that people who "do well" have access to some academic and learning "magic".

Once the student begins to know what the playing field of the discipline is, they have some awareness of what is required of them should they "have a go", and thereby give themselves a chance of success. Acceptance of where they are, where they need to go, and what they need to do to get there is not a painless or risk-free process. However, their willingness to have a go, and to begin approximating what is required is greatly enhanced if the student is able to risk translating or reframing the assignment in the context of their own language and their cultural and social context — their way of being in and making meaning of the world. Cambourne's "conditions for learning" have been utilised as a tool for reflecting on and making meaning of the student's learning process, practices, and needs, and for assessing the appropriateness of the teaching and learning process. In turn, this has been used so that the learner's process/practice and needs can be aligned with instruction and learning facilitation that meets student needs. The "conditions for learning" have also contributed to the mix in which tasks in the DTW process were specified (Gluck, Draisma, Fulcher, & Worthy, 2004).

Joining and Working through Kinaesthetic and Spatial Intelligences

The following example provides insight into working through a specific student's kinaesthetic and spatial intelligence to enhance their process of repositioning or bringing themselves into relationship with the culture of learning, language, and patterns of working within the discourse. The example also tells how using the student's dominant intelligences enabled them to reposition or redefine themselves as learners, and to use images and their body experiences to make meaning of the world. These images and experiences provided a starting point for the student to begin to tell her story of the assignment, firstly in her own language, and then in the language of the discipline. Students become multilingual as they begin to acquire and master the language and ways that work is carried out in the culture of the disciplines in which they study. As many indigenous students have exclaimed:

When I go home I get trouble because my mates tell me I don't speak right any more. They tell me I speak university. So I have to be careful when I go home to speak like they do. I allow myself to fall back into it. It becomes automatic after a while.

While at school, this student focused only during dance classes. Following school she entered classical dance education at university; however, a knee injury prevented her from graduating and pursuing a career as a professional dancer. It transpired that she now wanted to become an environmental scientist, but without possessing any of the requisite background in mathematics, biology, or chemistry: "I told you, I've been dancin'".

The student was brought into relationship with chemistry and maths through her kinaesthetic and visual strengths, together with the disciplines she had acquired as she was apprenticed into the thinking of the world of dance. The process of immersing the dancer in the disciplines of chemistry and mathematics so that she could receive demonstration of how learning and assessment were structured in the culture of these disciplines began with joining, listening, and dealing with her expectations of herself and her fears. For example, she had never studied chemistry nor written a university essay.

Her first assignment required her to consider the differences between gases and liquids. She was at a loss as to how to begin to respond to what was required. Despite the learning facilitator (Gluck) never having studied chemistry, he nevertheless had prior experience in enabling Aboriginal students to succeed in chemistry (Draisma, Gluck, Hancock, Kanitz, Price, Knell, Sharman, & Squires, 1994) and physics (Gluck & Draisma, 1997).

We began by listening to her fears and expectations of herself. It soon became clear that she needed to become aware of what she was required to do in order to have a chance of success. In effect, she required demonstration of the written form that her assignment needed to take; in other words, "What's it supposed to look like?" She also needed to know who to talk with, how and where to acquire relevant information, how to assemble and consider that information, and how to present her analysis (in other words, production).

Once the listening process was completed, the student's experience of dance was used to reframe the assignment. She needed to experience and come to know that the learning processes and skills acquired during her dance education were transferable and useable in science education.

The student experienced connection (transferability) as the office floor was cleared and used as a stage on which to assemble the attributes of gases and liquids. Student and facilitator then danced the attributes around in a manner that allowed the former to compare and contrast the nature of gases and liquids. Information on attributes was obtained by reference to class notes, texts, and asking other students (more capable peers) within the science disciplines. The student was then assisted to recall images of the choreography and dance process and utilise it as a basis for structuring, recording, and writing draft text to approximate what the chemistry subject required. As she produced assignment drafts, she was encouraged to reflect on the work process so that she was aware that each draft provided a building block upon which she could build her argument — her story of the assignment, as it were. As the drafts approximated what could be acceptable to the markers, the student was introduced to the idea of taking her draft paper to the lecturer to seek his expert commentary. The paper was then revised, submitted, and subsequently passed. The work provided a model that enabled her to begin to be apprenticed into a journey from being a dancer to becoming a writer.

The student was also able to utilise her understanding of the importance of drills and skills in learning and mastering the disciplines within dance to appreciate her need to develop drills and skills that would support her acquisition and mastery of chemistry and mathematics. For example, during mathematics classes she was able to understand what was being demonstrated, and to apply the principles within the demonstration to solve problems within the structure of the class and presence of the tutor. However, once she left the classroom and was presented with an assignment, an examination, or in-class quizzes, she was unable to repeat the problem-solving process. A meeting was called with an empathic maths tutor. After some discussion, the student and the tutor understood the importance of drills, skills, and practice in becoming consciously (and unconsciously) skilled in learning and performing ballet. It was then a small step to bring into her awareness the importance of drills and skills in performing and learning to apply mathematical skills and knowledge to a range of problems in a range of settings — quizzes, assignments and examinations. A series of mathematical drills and skills and practice problems were developed with the student. As she began to use them to reinforce the learning that had taken place in formal classes, she became unconsciously skilled in using the tools of mathematics and enculturated into its language and thinking.

The process of building approximations through discussion and images, and the development of drills and skills to enable her to go from consciously to unconsciously skilled, was repeated as the student went on to other assignments. After a few assignments, she required less and less assistance. The process was not smooth or emotion free; it was at times highly charged, and required the student to continue to make herself willing to ask for and use help from more expert others as she journeyed toward becoming and recognising herself as a writer.

The notion of travelling or journeying from where she was to where she needed to go, in other words, toward her potential, with the help of more expert others is akin to traversing Vygotsky's "zone of proximal development" (ZPD), "the distance between the *actual* developmental level as determined by independent problem solving, and the

level of *potential* development as determined through problem solving under adult guidance or in collaboration with more capable peers” (Vygotsky, 1978, p. 86).

The nature of the relationship between the dancer and the facilitator is encapsulated by Wells (1999) in the following assembly of extracts from Vygotsky’s *Thinking and Speech* (1937/1987):

leading (her) to carry out activities that force (her) to rise above (herself).
(p. 213)

The teacher, working with the (student) on a given question, explains, informs, inquires, corrects, and forces the (student herself) to explain. All this work on concepts, the entire process of their formation, is worked out by the (student) in collaboration with the adult in instruction. Now, (when the student) solves a problem ... (she) must make independent use of the results of that earlier collaboration. (pp. 215-216)

The process of facilitating dancing, recalling images, and recording her story into text, then using that as a basis for developing approximations and a final submission with the help of more expert others was not a straightforward or simple process. Rather, it was a means of evolving an experience and generating images and text that enabled the student to use her way of making meaning to construct her story of the assignment and bring it into alignment with the requirements of the discipline. As she risked dancing, recalling and using images, making approximations, and putting them into text so that the story was outside of herself, enabled her to stand back — to stand outside of herself — to see what she was creating and how it was being created. Moreover, these creations, movement, images, and text were part of her structuring knowledge and her becoming a writer. They were tangible, observable signposts of her process of becoming a writer.

The student does not remember the dancing that took place in the learning facilitator’s office. By contrast, it remains a strong memory for the latter, because it is a vivid reminder of the value of generating images and body experiences when facilitating learning with people who have dominance in kinaesthetic and spatial intelligence.

THE ‘ANTI-IDIOT’ SPRAY

The following story encapsulates some of the difficult times experienced by this same environmental science student during her journey. It also tells how the student’s being a dancer and performer provided a medium to work through those challenges.

A chance encounter outside the Aboriginal Education Centre produced the following exchange:

Russell: *How’s it going?*

Student: *All right.*

R: *What are you working on?*

S: *An essay for environmental studies.*

R: *When’s it due?*

S: *Next week.*

- R: *Where are you up to?*
 S: *I've completed it.*
 R: *Where is it?*
 S: *It's in my bag.*
 R: *Can I have a look?*
 S: *Yeah.*
 R: *Bring it down to my office and I'll have a look now if that's ok?*
 S: *When I've finished me smoke.*

Shortly afterwards, the student arrived with essay in hand.

- R: *Have you shown this to anyone else?*
 S: *What do you mean?*
 R: *Have you run it past your tutor or lecturer, you know the people who are going to mark it so that you can check whether you have covered what they want and put in a form they require?*
 S: *No way.*
 R: *Well, that is what consultation times are for, to get feed back so that you can redraft and submit before it is marked.*
 S: *No way! I'm not going to do that!*
 R: *Why not?*
 S: *Because I don't want to look like an idiot! If it is not up to scratch, I will find out when I get it back. It'll either pass or fail. And I am not going to be humiliated by the tutor or the lecturer. I am not going to feel or look like an idiot as they tell me my draft is no good and then tell me the final is no good. I would rather just hand it in. It either passes or fails.*

At this point, Russell pushes none too politely past the student and yells down the hall to the Student Counsellor located three doors away:

- R: *Hey Fred, have you got the Anti-Idiot Spray there?*
 Fred: (from inside his office) *How much do you need?*
 R: *I'm going to need a carton of the bloody stuff because we have a major case here. Have you got a minute?*

The student was leaning against the doorframe and looking for escape from the "hallway theatre" into which she was being thrust; however, there was no escape because both facilitator and counsellor were blocking her retreat.

- F: *What's happening?*
 R: *She has done the essay and doesn't want to show it to the tutor or lecturer before she submits. She's frightened if it's not up to scratch, she'll look like an idiot! She reckons she'd be better off submitting without feedback and redrafting.*
 F: *Yeah, she needs the spray bad.*
 R&F (in unison): *Getting feedback and using it to redraft your final submission is what it's about. That is the Anti-Idiot Spray.*

- S: *I'm tired of doing drafts and getting feedback. I just want to write the bloody thing and put it in. It shits me that I can't just write the bloody thing straight off!*
- R: *Welcome to the world of writing. It takes time to draft and redraft, and as you know it takes arse glue and a fist to write. That is all that is required, arse glue and a fist. You glue your arse in the seat, produce a draft, use your fist to knock on doors of people who know what is required, get their feedback, and then rework and submit. It's better to appear like an idiot at draft time than to prove you are one when you get a fail because you didn't follow the pattern that has proved successful for you.*
- S: *Yeah (with great resignation).*
- R: *So who is the tutor?*
- S: *Greg.*
- R: *Do you get on all right with Greg?*
- S: *Yeah he's good.*
- R: *So use my phone and set up an appointment.*
- S: *Ok.*

The student then phoned to arrange an appointment, went and met with the tutor, redrafted and submitted her assignment (and subsequently passed).

- R: *Just prior to her graduation and taking up employment in the environmental science field the student shared that her strongest memory was of the two of us walking down to the Faculty Office, me standing at the end of the hall way blocking her escape as she proceeded to the tutor's door, made a fist, knocked and asked the discipline's 'more expert other' for help and using it to redraft and submit — and the Anti-Idiot Spray continued working through the rest of her degree.*

Telephone conversations with the student after she took up employment have also revealed she continues to utilise more expert others as she continues her learning journeys through her professional life.

Stop “Shoulding” on Yourself

Difficulties that the students experienced with their writing are related to similar difficulties that the learning facilitator has likewise experienced (and continues to experience) when writing, together with what the latter (and others) do in practice to overcome them. These stories tell of fear of a blank page (or blank screen), feelings of inadequacy, and being physically, mentally, and emotionally taxed by the boring activity called writing.

- R: *I take great pleasure in telling the students how I failed English every year I was at school and how I regarded myself as 'illiterate or literacy inefficient', even after I had completed a number of university degrees. I relate how I felt I could never say what I wanted to in a way that I perceived got my message to the audience I was writing for. Most of all, I retell how I didn't meet my expectations of what I thought I should be able to do as a writer. For example, I retell the feelings that*

I had as assignments or refereed articles would be handed back — a feeling of doom. I also tell them stories of how I have recreated my story of writing so that I can now experience writing as enjoyable, exciting, and an adventure in telling my story — an adventure in which I no longer ‘should on’ myself. I relate and act out stories of how I was ‘should on’ from a great height by teachers, colleagues, and most of all myself. My self talk was ‘I should’ be able to write like other people — that mob that gets ‘top marks’ or has refereed articles easily accepted.

In sharing experiences of writing, it is emphasised that the writing process is personal — it is about *me*, it is about *my* story. Whatever the assignment or the writing task, it is acknowledged that the student is taking a risk by telling their story of the assignment or thesis. The process of writing is shifted from telling or reporting something objective and separate from themselves to something that is personal and that is theirs, something that they construct. It is also emphasised that their writing and their experience of writing is one of becoming and coming to know through building and telling their story.

R: I frequently act out in the theatre of my office: *Your initial draft may be excrement to you. And by the time you have finished, you will know your first terrifying death-defying act of Drawing, Talking, and Writing, and making the first draft contain the excrement that will give you the fertiliser you need to grow your work. Your fertilizer for growth starts with the basic excrement. So just get it down and stop ‘shoulding’ on yourself.*

JUST GET IT DOWN, MAKE A BEGINNING: STOP “SHOULDING” ON YOURSELF!

Students are encouraged to get down on the page whatever they think and feel is the story — to just make a beginning on the story they are trying to tell so that they can reflect on what they have put down: “Take what you like, use it as a building block for the rest of your story”. This emphasis invariably results in students exclaiming: “I *should* be able to sit down and just write what is up in my head, put it on the page, and just walk away. I *should* be able to do it all in one go!” The facilitator then utilises the idea of being “should on” to tell stories of how *they* were should on from a great height by teachers, and most of all by the latter’s unrealistic expectations of himself. The students invariably have no difficulty in identifying with stories of being “should on” by others, and most of all by themselves.

Once student and learning facilitator have joined through the “should ons”, once they have personalised the writing process, once they have shared stories, they are able to come to awareness and acceptance, and then begin to take actions that enable the development of another way, a way that enables them to tell the stories of what they want to write in a way that meets the needs of their audiences and themselves.

Just like the dancer (see previously), most students are initially resistant to the idea of personalising their story of their assignment, let alone draw, talk, and then write their assignment. However, there is nothing like desperation to generate willingness to join and try another way. By the time students have reached the facilitator’s door, they are

generally “running scared” and running out of options for staying in the institution and completing their studies. Consequently, they are generally prepared to roll the writing dice one more time. At some time in the joining process they are asked: “What is the definition of insanity?” They generally respond with incredulous expressions and silently consider running out the door. But when desperate enough students *do not* run, an opportunity presents itself to talk about the insanity of continuing to employ a process of writing that has *never* worked for them in the past, and expecting it will miraculously work in the future.

Stumbling on Another Way with Desperate People

The Koori Mail — an Aboriginal newspaper — was publishing students’ personal accounts of who they were, where they had come from, what they were doing before they had come to university, what they wanted to come to university for, what they did to get into university, and how they were being supported and succeeding in their studies. Aboriginal inmates in a maximum and high security jail read these stories and approached their education officer to ask if the learning facilitator (Gluck) would go to the jail and talk with them. The gist of the conversations that subsequently took place between the Head of the Aboriginal Education Centre, Gluck, and the prisoners could be summarised as follows:

Prisoners: *We read those stories in the Koori Mail and we need to get some of that literacy. Otherwise we are going to get out of here and find ourselves comin’ straight back inside. You know some of us have been doin’ life on installments — in and out, in and out the door. The main gate is like a bloody revolvin’ door.*

As we sat there trying to think how to respond, our eyes drifted around the cavernous brick room painted in prison green. Our eyes focused on the only decorations, a few paintings in the Aboriginal xray and dot painting styles of northern and central Australia.

We: *Who did those paintings?* (A few of the men identified themselves as the artists). *You like to paint?*

Spokesperson: *Yeah we like to paint and paintin’ isn’t gonna earn us a livin’ and keep us out of here [prison]. We’ve had paintin’ courses. We can paint. We need to get some literacy. And we don’t want that TAFE² mob, we’ve tried their literacy and it doesn’t do anything for us. It’s just bloody boring. Shit, we told em not to come back. Last thing they wanted to teach us gardening! In here!*

We replied: *Ok what about we teach you literacy through art.*

The spokesperson replied: *Yeah that sounds like a good idea.*

Accordingly, it was arranged to work with the prisoners to enable them to become literate through art. The prisoners did lots of painting but little writing. Some prisoners began to tell the story of their painting, and this was used in turn to build text (including a few poems); they felt safe enough to talk with the facilitator about their painting. An exhibition was subsequently held at the University Gallery, with the Regional Commander for Corrective Services opening the show.

Jail security protocols made the use of tape recorders and computers problematic. Consequently, when the prisoners were not able to write, or talk and write at the same time, the facilitator recorded what the prisoner said by way of handwritten notes, which were then used as a basis for further drafts (which were left in the hands of the prisoner for his own safekeeping). Both the painting and the story being told were the property of the prisoner; moreover, they were the experts on the story of their story. Consequently, they had control and ownership of what was painted, spoken, and transcribed into text, as well as the power to adjust, redraft, or delete whatever they had to say; they also had the power to participate, observe, or withdraw. Most significantly, whatever they said was correct from *their* perspective, since they were the expert; in a prison environment, this is a remarkable state!

Not only did the facilitator learn about people being expert on the story of their story, but also how to work with enabling them to tell and build that story, so they were able to use *their* language and *their* context (social, historical, and cultural) as the basis for being and becoming a student, and expert in learning. The process of eliciting their story and bringing outside resources and expertise to their process in a manner that enabled them to reinforce and value their expertise was critical to their repositioning themselves to repaint, retell, and rewrite their story of themselves. The value of their story and expertise was reinforced by listening, observing, and acknowledging what they had done, and then introducing material into the discussion in a manner that enabled them to access and utilise what they liked and discard what they did not. The key was not the resources, but rather it was the process of joining, safety, and the nature of the interpersonal relationship that enabled outside material to be brought into the relationship or learning space in a form that they could control and take what they liked and leave the rest.

In its most basic form, drawing, talking, and writing enabled the learner to risk presenting and telling in their language, that contained their voice, who they were, and to discuss it with the facilitator as both a “more expert other” and “a co-learner”. Where the prisoner considered that what was introduced to the discussion enabled them to enhance their story, they were then in a position to consider incorporating that material into their story and owning it. For example, a young man doing his first time “inside” came into the room and sat with the other students and watched them paint. After a short time, he picked up a sheet of paper and painted a kangaroo in the x-ray style of the Northern Territory. The fellow next to him was painting in the dot style of central Australia:

R: *Hi I'm Russell. (pause) That's a terrific kangaroo.*

Inmate: *Yeah (silence) Yeah. I'm painting a kangaroo. (long silence) What's this class all about anyway?*

A discussion ensued as to how the class had come about as a direct response to inmates' initiative. Once the politics of who we were and why we were there had been discussed, the inmate felt confident enough to say:

Inmate: *I've known for years I've got Aboriginal blood in me. But I don't know much about who I am, where I come from, who my family is.*

This information provided a massive amount of material that was able to be introduced into discussion with the student. Moreover, discussions in such a prison setting were not private; everyone else in the room was listening in and watching this interaction. Thus information about and themes inherent in his painting were able to be introduced in future discussions relevant to large numbers of the inmates.

The preceding short statement provided direction for the introduction of material on the connection between painting and law — painting tells a story, what the symbols mean, what country they relate to, ownership of information, racism, and arguments of blood and identity. As a direct result of this brief discussion, materials (videos and books) that provided information on the topics were able to be located. Furthermore, and with the prisoners' agreement, genealogical and historical records were accessed at the Institute of Aboriginal and Islander Affairs, in order to begin to determine who their families were, where they came from, and what piece of land was their country. The process and the introduction of this information led to a difference in the way the men talked, painted, and (in the case of a few) wrote. As they assimilated the information, their discussion, drawing, and painting presented a story of what their work was and what it was becoming. They also presented a story of themselves as someone who was becoming a learner, a storyteller, and a writer.

THE LANGUAGE OF PRESENTING OURSELVES

The differences in experiences and the language used to present ourselves to each other provided a rich opportunity for all to learn. A shared understanding of language gradually developed through working together. As intimacy of language grew, this enabled the integration of information from outside each other's experience into individual stories. This is equivalent to Lysaght's (2001) emphasis on the importance of language and the development of shared language in the learning and development of female adult students from non-traditional backgrounds "We present ourselves to the world through the language we use and, to a large extent, our own understandings are developed through the access we have to the languages of others" (Vygotsky, 1962, 1978).

DTW provides an arena in which to share stories and experiences, to the extent whereby a learning facilitator could become a co-learner in a high security prison. Conversely, prisoners could present themselves as not only prisoners, but alternatively as painters, discussants, learners, writers, experts, and custodians of their story of themselves and what they were becoming. DTW gave prisoners a process, a tool for portraying who they were, standing outside themselves/their story that allowed them to observe, reflect, and discuss themselves and their reflections with themselves and, if they chose, with others.

The adoption of DTW as a tool for individual and group reflection was evident in one of the men's work. He was a prolific painter, and most of his painting was done in his cell, alone; he painted little if anything in class and was continually seeking information on culture and prisoners' rights. He also started the idea of a Koori prisoners' newsletter, that led to conversation on how the artwork and what was written could be

shown to the outside community through the Koori Mail (as it eventuated, the idea was not able to be pursued because he was released prior to the arrangement being made).

The idea of the Koori prisoners' newsletter enabled the prisoners to begin to construct, tell, or portray their own stories to themselves, to be their own audience and, if they wished, allow others to be members of their audience, an audience with whom they could choose to share their stories and receive feedback. They had control of what to let out, what to let in, and what to keep private and unchanged. Most of all, it provided them with the *safety* to tell their own story to themselves and witness it.

DTW provides a process in which it is possible for the prisoner to become aware of their story, to accept it, and be apprenticed to being and acting on their own behalf by using their "dark side" to write a new positive story. Was the prison work a success? Firstly, a few prisoners upon release from high security made their way to the University and spent a few weeks working at their own pace and in their own way while they thought about coming to university officially and what they wanted to do. All of these people had addiction issues and were subsequently linked with rehabilitation programs. Secondly, one postgraduate student claims the DTW process allowed him to bring himself into relationship with his past, his worst stories and fears that were depressing and preventing him from releasing himself to write: "I needed the safe place, a means of getting there, and the ability to recognise that I had to give it to myself. Drawing, talking, and writing allowed me the safety to risk taking the journey".

MORE DESPERATE PEOPLE AND THE CONNECTION TO COMPUTERS

Work with indigenous prevocational and vocational students attending TAFE or incarcerated in prisons and university undergraduate and postgraduate students has led to an appreciation of the role of drawing, talking, and writing in people's journeys to literacy efficiency. Working with Aborigines in a range of educational settings has provided experiences that led to the idea that computer software could be produced that incorporated drawing, talking, voice recognition, and writing. Ideally, the software would respond to the student rather than demanding that the user adapt to the demands (constraints) of the machine.

A period of study leave was undertaken by Gluck in the spring of 2000, working with indigenous visual arts students at an Aboriginal Technical And Further Education — TAFE — college. The majority of these students were literacy-inefficient in the context of their course work. In addition, some could not read the captions under the picture of the sports star on the back page of newspapers. By the end of this study leave, some of these students achieved literacy efficiency that would have allowed them to enter, and with help successfully complete, undergraduate studies at university. Moreover, the students had become computer literate, and were able to utilize a computer as an integral part of their DTW process.

MAKING SENSE OF THE STUDY LEAVE

At the completion of this study leave period, three students and a teacher were asked: “What’s it been like for you working with me?” Each student was approached separately, and their account of the experience recorded, typed into text, and then given back to them for checking and amendment. Each student was then given a copy of their completed transcript and a copy of the interview tape. It was agreed that the facilitator could keep a copy of the transcript and the accompanying tape for study purposes. A similar process was undertaken with the teacher with whom Gluck had worked very closely during his time at the college.

The interviews were undertaken at the end of the college year, a year that resulted in students developing and using their stories to become literacy-efficient visual artists in the creative arts discipline. The students’ processes culminated in them being able to risk collaboration with music and theatre students, so that elements of their stories were performed and presented in public to the arts community, their families, friends, and colleagues. The full text of all the stories that the visual art students had developed during the year was displayed in books on the walls of the performance space.

The transcription of the students’ and teacher’s interviews took place during the college summer vacation. The process of typing, checking, and amending the transcripts yielded riches beyond words. The process of transcribing the interviews with the teacher and two of the students was relatively straightforward. The patterns or combination of words and meaning were easily identified and transcribed. However, the transcription of the third student’s interview was exhausting — his 90-minute tape took many weeks to transcribe. Listening to his tape and attempting to transcribe verbatim led to repeated loss of focus of the words spoken and comprehension of what he talked about. Eventually, it was realised that he had filled his speech with thinking and cultural safety phrases such as “sort of thing like that”. For example, his speech went something like this:

I was walkin’ down the mission, sort of thing like that with my mates sort of thing like that when we ran into the mission superintendent sort of thing like that and he sort of thing like that asked us sort of thing like that where we sort of thing like that were going sort of thing like that. We stopped sort of thing like that and said we were sort of thing like that going sort of thing like that to the river sort of thing like that to fish for red fin.

After recognising the patterns in the speech, it was possible to go back to his tape and easily make meaning and record verbatim. The process was identical to working one-to-one with the student on the computer: he talked, and the learning facilitator automatically filtered out his thinking and safety phrases. When these “phrases” were edited out, the following text emerged:

I was walkin’ down the mission with my mates when we ran into the mission superintendent. He asked where we were going. We stopped and said we were going to the river to fish for red fin.

However as Ernie Blackmore, an indigenous PhD student in English, and familiar with DTW, observed:

And this [the version where the phrases are edited out] is crap! Although it tells the 'story' it is not the 'voice' or the intent of the narrator. It is homogenised beyond the point of recognition.

Safety phrases remain an integral part of the student's everyday speech and urban voice. Everything that he said, including the thinking and safety phrases, was re-recorded into text. When this transcript was presented to the student at the beginning of the next school year, his use of safety and thinking phrases was highlighted using strike-throughs:

I was walkin' down the mission, ~~sort of thing like that~~ with my mates ~~sort of thing like that~~ when we ran into the mission superintendent ~~sort of thing like that~~ and he ~~sort of thing like that~~ asked us ~~sort of thing like that~~ where we ~~sort of thing like that~~ were going sort of thing like that. We stopped ~~sort of thing like that~~ and said we were ~~sort of thing like that~~ going ~~sort of thing like that~~ to the river ~~sort of thing like that~~ to fish for red fin.

The student immediately recognised the safety and thinking patterns in his speech and recalled what he was doing/thinking when he used them to deal with the management arm of the mission and bureaucracy in general. He also recognised that these patterns continue to be an integral part of both his everyday speech and for the process of transacting business. They have become an integral part of his urban indigenous voice as he negotiates with education, social security, housing, medical, and a range of other bureaucracies and social and cultural contexts.

Upon witnessing his use of cultural and safety "phrases", the student experienced an "ah ha", Eureka, or "light bulb" moment. He was able to see how his speech could be put into text and easily edited so that it was readily understandable by others. He recognised that if the computer could put his speech into text, safety and thinking phrases could either be suppressed or noted with a strike-through. This would give him a choice of texts. The computer was a tool for mediating and communicating meaning across cultural contexts.

The strike-through process also provided a means of demonstrating to his teachers that he was not illiterate or literacy-inefficient. Rather, it allowed teachers to recognise his dexterity and inventiveness when dealing with and working across complex social and cultural domains/context. From this point on, teachers began to see him in a different light. This was not a person who could not speak a coherent phrase; this was a highly accomplished reader of cultural and social contexts who could develop and create stories and communicate them in a manner that was appropriate to ensuring his safety. Once the student and his teachers had gained insight into his way of voicing and communicating his story, he began to experiment by risking to build and communicate intimate stories and arguments that were suitable for a range of audiences. For example, as the *strike-through* section of transcript was presented to the student, he was able to talk about and

provide other layers of information that detailed the place of social and cultural safety phrases.

The discussion about the phrases then provided an opportunity for the storyteller to provide insights to his process of thinking, acting, and being:

***I was walkin' down the mission,** ~~sort of thing like that~~ [thinking pause]
with my mates ~~sort of thing like that~~ [thinking pause feeling his way into
the context of the mission]
when we ran into the mission superintendent ~~sort of thing like that~~
[thinking pause continuing to feel his way into the context of the mission
and beginning of a safety phrase]
and he ~~sort of thing like that~~ [furthering safety phrase — how much is it safe
to divulge to the listener/teacher/transcriber]
asked us ~~sort of thing like that~~ [safety phrase with respect to the listener and
the frame of reference — the superintendent]
where we ~~sort of thing like that~~ **were going** ~~sort of thing like that.~~ [safety
pause and contexting continued]
We stopped ~~sort of thing like that~~ [thinking and safety phrase and then
decision to provide the next level of information]
We looked at each other, [visually] for hints on what to say [as we built a
safe response to the superintendent] because he had the power to cause us
all sorts of problems. [In those days the superintendent had the power to
withhold rations, determine who could stay on the mission, and much
more.] You had to be careful what you said. So we would let little bits of
information out and see the superintendent's reactions and then add some
more that we thought was safe. We would also be looking into each others
faces [and body language/signals] to check out that we were giving out safe
information. Every one of us walking down the mission knew talking with
the superintendent could be dangerous [for the people talking and for
others referred to.]
and said we were ~~sort of thing like that~~ **going** ~~sort of thing like that~~ **to the**
river ~~sort of thing like that~~ **to fish for red fin.***

Now this extra level of information could be built into the story if the teller wished. The computer makes it easy to incorporate this extra level of information providing you have a safe environment in which someone can utilise a tool to record and put the conversation into text form. The ability and facility to converse, to narrate, tell and record the story with the tip of one's tongue rather than through one's finger tips allows the teller to incorporate safety phrases into the text, and then to decide what to do with them. Decisions on what to do would be influenced by many factors, such as the cultural and social origin and physical location of both teller and listener, and the sense of audience of both.

PUTTING STUDENTS STORIES INTO TEXT ON SCREEN AND READING IT BACK: A TOOL FOR ENHANCING LITERACY

Work with one of the more literacy-inefficient students started with the learning facilitator acting as a word recogniser; the student talked, and the facilitator typed and read back what the former had said. The process of reading the student's transcribed story back to him as he followed the text on the screen often resulted in their intervening and providing more information, which was subsequently added to the story. More importantly, it was noticed that his lips began to move in sync with the facilitator's voice as the latter read the story back to him from the screen. The student had begun to recognise the patterns of words and phrases in the text of his story on the screen. Then one magical day he said:

You didn't read exactly what's on the screen. You made a mistake.

At this point it was clear he felt safe to begin to transform his place in the learning relationship from student toward co-learner. Not long after this event, the student began to regularly correct the recording and reading process, by saying: "You didn't record or read exactly what I said". Shortly after this, the student began to read aloud what was on the screen and began to use the keyboard to edit his story.

The student's process of mentally (silently) reading back his own material promoted his acquisition and control of text and language. Even more importantly, it was discovered that his reading back of the text reflected his anticipation of what should or could be on the screen, a combination of memory, anticipation, and meaning-making. Significant divergence between what he anticipated he had said, what was recorded, and what was read back often signalled that information was missing from the story. This frequently led the student to develop and include a whole new level of information and direction that enhanced his story.

Having access to his own story on the screen facilitated the student's preparedness to risk reading aloud, because he knew the context and detail of the information in the story intimately. Consequently, he could anticipate what was going to be on the screen. The idea of "mistakes", together with the courage it takes to experiment, was often discussed with the student. In discussing mistakes, it was emphasised that the story was *his* story, *he* was the expert of his story; nobody else knew or owned his story. He had the power to tell and control what he told and to whom, in other words, who had the right to access his story. The computer provided a good deal of safety because the student was able to put his story on a disc, take it out of the machine, and put it in a safe place. This can still result in withholding all the "truth" from disclosure because the process does not always guarantee safety.

The process of reading the text back to the student as he followed it on the screen often gave him the space to listen to and evaluate whether the information he was divulging was appropriate for the audience it was to be delivered to. It also gave him the facility to question and be questioned whether the story flowed and "hung together".

- Russell:* I well remember the day he was reading back, and it became obvious to me that his text did not flow. The story was flowing, then it jarred, kicked, bucked just like a chain saw hitting a nail in a log. He quickly moved on, and at the end of his reading, I read his passage back to him. I used my voice and actions to suggest something was missing — a relationship hidden, not divulged. At this point we were standing face to face, at most 15cm between us; he verbally exploded as I said:
- R:* This doesn't work. Tell me more about this relationship and how it connects to this. *(There was some unfinished business; we were talking about a very close relative that was mentioned in his story.)*
- S:* This is getting too personal mate. I'm telling you mate! You mate are not going to know this stuff!! This is my stuff!!! Family business! *(emphasis on "mate"; vehemence with which the words were spat; I was covered in spittle).* We stood staring at each other, which culturally is a huge confrontation and can proceed to getting the crap beaten out of one. I stood still and said:
- R:* Writing is personal.
- S:* This is too personal this is my business!
- R:* I continued to stay very still. I dropped the level of my voice to a normal speaking voice and we began to be able to relax the process. And I proceeded to say: Well it's your decision as to what goes in here.

The riches that flowed from this reading back process were huge. From this point on, the writer recognised the value of reading back, the use of voice, body, text, and visuals. The floodgates opened, and he began to write like he had never written before. He came to know he was writing about himself *for* himself, in other words, for self-exploration and growth. The question of safety and security of information still remains, however.

The learning facilitation process was far removed from repetitive drills and skills. The student experienced the computer as a tool that he could safely experiment with and use to develop, to control, to tell and communicate his stories while simultaneously learning to be computer-literate and literacy-efficient. For example, he was able to make statements to the effect of:

The computer is a tool. It allows me take the boring crap out of writing. Like when I was at school it was exhausting [the process of drafting and redrafting with a pencil was physically and emotionally taxing]. Bloody boring [tedium, the repetition] and having to get it right! [the teacher's way, the way of the mainstream in a language and context that was not relevant to the student's private, cultural and social world. The world beyond the incarceration space of the classroom]. Anyway I was put down the back of the classroom next to the window so that I could stare out the window: 'Be quiet! And just wait for class to be finished'.

The student later went on to research and write his own material and began to combine his visual, oral, and textual abilities with his telling stories of his people's survival and growth.

By using the student's context, stories, and the language of his labour of everyday meaning-making, he had a rich base for facilitating his acquisition and control of Creative

Arts language and learning. I remember him telling me one Friday afternoon how his writing was informing his artwork. A base had been laid that opened the possibility for him to further his study within the Creative Arts disciplines and enhance his creative process through writing. The process of becoming literacy-efficient through his own stories also provided a starting point for thinking about what form a software implementation of DTW could take. Safety issues still remain paramount. A way through the safety process is enabling students to tell their stories and fictionalise reality, in which we switch the truth for fiction and fiction for truth, and create a place in which they can work safely.

First Adventures with Speech Recognition: Smashing Affairs and Serendipity

Following the aforementioned period of study leave, experimentation commenced with commercial speech recognition software, namely IBM ViaVoice® and Dragon NaturallySpeaking™. More specifically, the learning facilitator read the text output displayed on the computer screen back to the text-illiterate user, after which the student repeated the phrases, in other words, *manually* closing of the feedback loop. Work also continued one day per week with the vocational students.

Any idea of using commercial, off-the-shelf speech recognition software was abandoned after a number of near “smashing” experiences with indigenous university and TAFE students. For example, the students’ frustration with commercial word processing and speech recognition packages boiled over to the point where some students physically attacked the computers!

It soon became apparent that word recognition software needed to be developed that was compatible with the needs of literacy-inefficient, orally-proficient, visually-strong computer illiterate people from an oral story-telling tradition. More specifically, feedback ideally should be provided in the student’s own voice, and not the learning facilitator’s (or even worse, synthetic voice output as produced by speech recognition packages).

At the same time, work resumed with Marion, an undergraduate student in early childhood education, as she worked toward detailing a rationale that supported her literacy learning facilitation practice. She was encouraged to pose a learning situation that she may be faced with when she was working in her community. Next, she was asked to think of the “worst case” classroom scenario she could imagine, tell her story of the scenario in her own language, and then put it into text (the language of which was subsequently refined in order to be acceptable to other academics in the faculty). She was then asked to come up with a strategy for solving this “worst case” scenario, and to detail the relevant theorists and their theories that underpinned her proposed practice.

Marion then proceeded to generate a narrative (a dream sequence) in which she was visited by the theorists, and asked them questions on what they would do and how they would utilise their theories to facilitate the child’s learning. Within her narrative, the theorists asked her questions and demonstrated in dance how they would intermingle their theories and actions to meet the child’s learning needs. Marion’s narrative was recorded, refined, and used as a basis for her to begin to explicate her rationale for teaching: “....dialoguing with real or imagined others ...is an essential part of the process

of textual composition that even the most knowledgeable others are able to continue to learn in the Zone of Proximal Development” (Wells, 1999).

We realised that the work was in a very different genre than what the marker was accustomed to receiving. The work was submitted and with a small amount of representation the case for a different genre was accepted and the essay was passed. Theory and the language of the discipline were being recast as a tool for her use. DTW provided a process for the student to develop a coherent whole between context, practice and theory. It enabled her to write a story that could provide a new direction for her learning and teaching.

However when Marion came to discuss further assignments, she began her process from her position of developing and presenting arguments in ways and in the genre that had previously failed her. After much thrashing about, she invariably returned and continued to return to and use the drawing, talking, and writing process that underscored the initial success. It was frustrating, though, to see Marion repeatedly start her investigation and development of assignments from a position that, despite being previously accepted by the faculty, was *not* one that served *her* learning needs. A subsequent altercation with the learning facilitator led to the explication of the process and a first draft of the steps involved in “Opening the Door to Literacy”. These were subsequently incorporated into a submission to the Apple University Development Fund, through which scheme initial funding for the development of voice recognition software based on DTW was forthcoming.

The learning theories that provide the theoretical underpinnings of DTW were presented at appropriate conferences (Fulcher, Gluck, Worthy, Draisma, & Vialle, 2003; Gluck et al., 2004); indeed, the coming together of researchers from diverse fields could be viewed as the equivalent of a *group* zone of proximal development and quasi-community of practice (not to mention a realization that similar task sets to facilitate literacy efficiency with students who were from different social and cultural backgrounds had been independently evolved and used by different collaborators). This coming together also allowed us to articulate the four related tasks that are at the heart of the draw-talk-write process, namely:

1. talking with the learning facilitator about the assignment topic and relating it to their personal experience,
2. utilising that personal experience to draw, decipher, and tell a detailed story or create a scenario that can be used to fulfil the requirements of the assignment topic or area of inquiry,
3. researching the topic and determining a theoretical perspective that supports the story as an example suitable for the assignment topic, and
4. refining the assignment through revisiting steps (1) through (3).

The above four-stage DTW process has been validated by way of “show cause” (restricted) students enrolled in the Bachelor of Teaching program at the University of Wollongong (Fulcher et al., 2003).

Enabling PhD Students to Develop and Write Their Dissertations

The most recent learning facilitation work has been undertaken with doctoral students, namely Ernie Blackmore (English, University of Wollongong), Frances Laneyrie (Management, University of Wollongong), and Robyn Thompson (Education, University of Canberra). The work takes place in the facilitator's office, which in reality operates as a theatre. An electronic whiteboard occupies a wall and serves as a theatre scrim upon which to talk out our thoughts and draw them on the board. This board is not state of the art; it does not connect to a computer, and it does not record voices. Consequently, sessions are recorded, and students take both tape recording and type with them when they leave. It is not always quiet and demure in this theatre/office; it gets intense and stormy at times (receptionists camped adjacent to the office sometimes object to the work because they cannot hear their phone conversations).

The tape recordings, scrim drawings, and the physical movement and interaction between student and facilitator all contribute to the students' telling of the stories that ultimately become their PhD theses; Ernie has coined the term "literary dramaturgy" for describing this process. Robyn experienced the process in a theatre that heavily relied on butcher's paper pinned up to the balcony of her house as she brought her thesis to completion. Russell's work processes were utilised by Robyn as a single site case study for her thesis. Others have experienced it as we have drawn with sticks in the red hot earth of north east Arnhem Land and the north west of the Kimberly.

The process does not require high levels of technology if you have people who can read and write and transcribe what is said or acted out into text in a timely fashion. However, the participants need to be really dedicated to the task because of the sheer physical demands and the time required for transcription. Ernie and Frances have a collection of tapes that they have yet to transcribe, but more importantly they retain images generated as a result of the experience in the theatre of Russell's office. The images are dated and provide a log of (access to) the conversations which transpired. Also, the drawings and their development act as much more than simply providing a log or a reflective space. As Ernie says, "the process is diminished when we do not have immediate access to the text — text that is synchronised with our voices, not a machine-synthesised voice, what's more."

The white board process enables us to work quickly and spontaneously and to take risks with thoughts that we could not do if sitting and writing. Writing is slow and laborious and frequently requires "arse glue" isolation and demands that we express our thoughts and what we want to say through the tips of our fingers on a key board, pen, or pencil; moreover, it requires us to be stationary, fixed in front of a computer or book in which we record. Alternatively, the drawing, talking, and writing process allows us to physically move and to use and be aware of our voice levels, body postures, and language, as well as the way we position ourselves in relation to each other and to "the scrim", within the theatre of our work. Awareness of our positioning in relation to each other, the scrim/whiteboard, and the forecourt or stage of our working theatre (the office floor) allows us to read each other's body language and emotional responses. It allows us to observe and practice protocols of challenge without confrontation and aggression, for emotional violence to be contained, and for superior/inferior positioning to be handled in a way that safety and generative relationships are enhanced. There is no place

for loss of face because we are able to flow and become used to evolving rather than forming fixed positions. Fixity is not entertained, and direct eye contact and the challenge of posture and gesture is used within cross-cultural bounds. There is no escape as we are focused in a way that reorients the learning environment to conditions for learning focused on the student and the facilitator acquiring language and making meaning by co-production.

The theatre of our work and, in particular, the scrim or screen and the forecourt of our theatre upon which we place, stir, boil, construct, and project our work allows us to record and layer the evolution of the argument. Printouts allow us to locate where our diversions fit in the overall of the pastiche. This visual location and tracking of argument provides a good deal of security for risking our diving into and developing a part. The doing allows us to know how to locate and shift as we develop the whole; it definitely enlivens the process. The sharing of the pen and the eraser greatly enhances the collaborative role — one pen user at a time. The person with the pen has the talking and painting stick. Each person is not independent; we are in relationship with the other in the drawing-talking-writing of this white board theatre. Each person realizes that what is unfolding or evolving is mutually dependent on the other, in other words, a co-production or collaborative enterprise. It fits with a traditional process of developing, teaching, and sharing knowledge. No one holds the key to “the knowledge” or to knowledge production. We focus on co-creation, interdependence, and knowledge dissemination similar to the theatre of ceremony where bark paintings are developed and utilized in the conduct of indigenous ceremonies. Confrontation, conflict, and creation are contained within the theatre of healthy mutual inter-dependence and co-production.

The white board theatre process allows exploration to be concretized, pictured, and reworked at speed because the stored image provides a record that allows the fear of losing or forgetting or being distracted to be accommodated.

The process is particularly suited to a thesis because it is about the who, how, what, where, when, with whom, costs, and benefits of urban indigenous voices, and the role of contemporary theatre in the development and conduct of these voices.

Ernie: There have been any number of walks around the campus, for whatever reason, and many of these walks have been an extension of the work in the “theatre” space that is your office. And on many occasions, although not as many as ought to have occurred, the tapes continued to role. Under these circumstances, other localities became the theatre, including many visits for lunch to Food Re-Thought, where we used their paper tablecloths as substitute whiteboards. It is interesting to note that when we returned to the “theatre” of your office the discussions continued with the white board being very quickly drawn back into the process, as it allowed for the arguments to be presented to and into a “safe”, familiar, trustworthy, but most of all a familiar place where there was a feeling that judgement was on hold. This sense of “freedom” then permits the free flow of ideas and information from both student and facilitator.

Out beyond ideas of right doing and wrong doing there is a field. I'll meet you there. Rumi (A Great Wagon, c1250).

The movement and the confrontation that can take place through the theatre or the “literary dramaturgy” should not be underestimated. It invokes movement by allowing all participants to collaboratively utilise the many different conceptions or ideas and fields that are encrypted and present in the fragments of their individual and collective drawing, gesture, and speech. Most of all, it allows the participants to explore and depict the relationship between what is and what is becoming; it enables an understanding of relationships.

Frances: *The drawing can be a lengthy process. It is not about getting the picture right; it is about defining and putting the parts down in an image and then getting them ... or drawing the relationship between them so that the relationship is right for me. Spending all the time getting the drawing or the relationships between the elements of the image right allows the writing to flow because I have the structure and the relationship/interrelationship sitting not just before my eyes — I have it in my mind and my body. My body remembers how I put it together; it is also the physical process and the relationship process with the physical layout of Russell’s totally shambolic office, which almost comes up to the state of mirroring my private working space at home. I can see the positioning of ourselves before the electronic white board, and I can also associate the emotional state that is present as I produce the drawings and depict the relationships. The emotional content is that I identify with my thoughts as interviewer and empathize with the interviewee in my study. Synthesis of my identification and empathy produced moments of awareness, clarity and acceptance that led to the action of drawing. Having gone through a similar sense of awareness to that being shared by the interviewees (in my study) and have that reflected back to me provided the light bulb moment As the drawing and the talking led to the awareness at the forefront of my mind, that awareness was drawn and accepted. Drawing in the theatre provides a mechanism or task for bringing seemingly unconnected ideas, concepts, theories, and experiences into a visual image, and then I am able to talk to the image, record it, and then use that as a basis for writing. The unpacking of the meaning and the connections in the image provides for a well-articulated space, an interactive space, because it allows for a conversation and a shared drawing space that leads to the production and knowing of something that would otherwise not have been generated.*

It should be pointed out that the drawings referred to above are artefacts and texts. They are a concrete expression, the embodiment of a principle, an abstract idea. The drawings are part of a semiotic mediation process that is housed in the literary-dramaturgy (the semiotic mediation process also applies to movement and dance). The drawings are even more powerful when read and used as a thinking device and a basis for building further text and meaning through dialogue with self and others. Importantly, the drawings are representative of what the drawer knew at the time of drawing, and are tools for creating a platform for further coming to know and extending the upper bounds of the drawers’ ZPD. The drawings made in the literary-dramaturgy process are also part of a co-production.

It is no coincidence that the emphasis throughout is on the importance of all writing (all research) being about story, a developing story, a story of becoming. It is the belief

of the authors that life is a story, and that the tasks involved in the process that underlies the DTW software in many ways provides a technology that facilitates narrative development and narrative therapy processes. It also parallels language development processes in children and adults, particularly the way in which private speech or internal speech is used, and the impact that voicing or putting thought into words changes the mental processes or development.

In Ernie's experience, the literary dramaturgy process provides him with the safety to write:

Many attempts have been made (to write), but if you're not safe you're not safe, and the resources are not resources. You have to join at the heart and the spirit and the soul. Safety is essential to enabling people to 'go within in order to go without' and 'to go without in order to go within'; safety to go into the secret space within themselves so that they can begin to prepare their story, the private space (to) access and assemble the story. Safety provides the means of going through the (seemingly) impenetrable barrier to the theatre of work that houses the potential.

...the accepted norm is the written word, black print on a white page. ... We do allow ourselves as adults to use visual aids, photographs, paintings, movies, videos, electronically and print-generated images to convey messages, Cartoonists tell stories in pictures without words, ... I can create the image and then use my verbal account to develop the text.

There have been times when Ernie has been absolutely stuck, unable to begin work. At this point, the learning facilitator has drawn a story or put an image on the board and proceeded to use the image to "spark up" or "kick start" the process. The following saying (aphorism) appears on the top of Russell's white board that evokes a picture that fits with the narrative process:

In order to go within you have to go without; In order to go without you have to go within.

Serendipity Continued

A 2003 Apple University Development Grant led to the authors working collaboratively with a graduate student and final-year undergraduate software project groups within the School of IT and Computer Science on DTW software development. The collaboration required specification of what functions the software was required to perform, as well as how it would be accessed and controlled by the client users (students). The SITACS students were initially keen to begin work. After an introductory meeting, regular project meetings were held between the students and the learning facilitator.

Meetings took place either in the facilitator's office or in a meeting room where one of the walls was taken up by large white boards. From the time of the first meeting, the facilitator began to tell and draw stories of how he had worked, facilitated, and been a co-learner with literacy-inefficient students as they journeyed towards literacy efficiency. Similarly, the Facilitator and students set out on their own journey to co-evolve the basis of the software for the DTW process.

The SITACS students appeared to listen to the stories from the perspective of identifying and developing components that would enable them to develop the human computer interface (HCI). They were intent on eliciting things to be done and aligning them with things they knew they could easily achieve. During the meetings, they wanted and continually sought directions/specifications. Instead, they were told stories of working with computer and literacy-inefficient people and the latter's journeys to literacy efficiency, which in turn generated ideas underlying the demand for this software. They contained thick descriptions of the relationships that developed during the journeys between learners, facilitator, co-learners, and technology. For example, the stories contained information on who the user population were, how they had successfully been enabled to access and utilise computers to enhance their learning, their attitudes to access, use, and security of information, what they used information technology for, and the contexts in which they were likely to use a DTW program. From the perspective of the learning facilitator, the stories contained critical information for the design and development of an effective HCI between literacy-inefficient people and the proposed word recognition software.

The students became irritated with the focus on story and relationship orientation. A major irritant for them was apparently their lack of experience in critically listening to stories and using the rich data within them to inform software evolution. This presented a challenge because one of the facilitator's major tools for conveying information is to tell stories that focus on relationships. Consequently, Marion was brought together with the SITACS students in order to share individual (and joint) experiences of the DTW process. Marion was able to contribute her experience of becoming literacy efficient through DTW, and the SITACS students were able to experience stories via Marion. This was particularly useful because in addition to her experience of becoming literacy efficient through her participation in the evolution of the DTW process, she:

- had contributed to the learning relationships that enabled the learning facilitator to begin to imagine, speak of, and define a direction for creating DTW software;
- is an active member of her indigenous community, a community that has a potential need for processes and resources that can facilitate journeys into literacy efficiency through drawing, talking, and writing;
- is an active member of a learning community within her home community, a qualified pre-school teacher, and literacy teacher, and is soon to be a qualified primary school teacher;
- is computer literate, well beyond the learning facilitator's experience, and has an intimate understanding of the technological milieu that members of the user population are exposed to on a daily basis; and
- is similar in age to the students.

Marion's incorporation into the work process enabled a number of tough situations to be mediated. For example, the SITACS students were concerned about their ability to integrate a drawing function into the software. Their response could be paraphrased as:

Our experience of becoming literate did not involve drawing; so it's not necessary to draw to become literate. Incorporating the drawing function

in the software is difficult or beyond our programming experience, so let's put it on the back burner. It's not a priority.

Marion was able to relate her experience of the value of drawing in orally- and visually- strong people's journeys into literacy efficiency. Consequently, the integration of the drawing function into the software came off the backburner. In essence, she played a mediating role between the information within the facilitator's stories, her literacy journey, her journeys with others, and the product-focused world of the SITACS students.

Her introduction into the work process contributed to the development of a shared way of working and language that enabled us to move beyond each of our individual stories that we brought to the project. We were then able to begin to produce something that *individually* we would not have been able to consider. We had entered a well-articulated space, we had a new way of talking and relating with each other and the technology, and we were becoming co-learners.

In summary:

- we were able to put our ideas into language and drawings that we could use as a basis for questioning and mediating meaning and developing shared direction,
- the students were able to be immersed in thick descriptions of how a computer and literacy-inefficient person who is orally- and visually-strong can utilise drawing, talking, and writing processes to become literacy efficient,
- the students began to consider and pursue possibilities beyond what they knew they could do easily and quickly, in other words, the familiar, and
- our collective direction was to inch toward producing a human computer interface that would go beyond our individual experiences of computer users and literacy journeys.

We had generated a relationship space that could be characterised as a well-articulated space, a space that enabled us to do and speak of what we could not otherwise have done (Latour, 1991). And within that space we operated a *group* zone of proximal development as we facilitated each other's learning. As Thompson (2003, p. 161) notes:

Bringing people into 'well articulated' spaces will allow a new language to evolve ... the importance of facilitating and engaging people in well articulated spaces, of bringing people and things into a relationship to enable them to speak in a different way. I am now compelled to focus on relationships rather than things, and the mutuality of those relationships rather than concentrating on 'components of a product' (Williams, 1997, p. 48) and objects per se.

It was in this spirit of facilitating a well-articulated space that Marion, the facilitator, and the SITACS students were brought together so that we could open ourselves up to possibilities. The process was fascinating from the facilitator's point of view and at times exasperating for the students.

The environment within the theatre of our work affords us the luxury of being able to risk aligning and defining ourselves with the unknown and what could be. Further, it

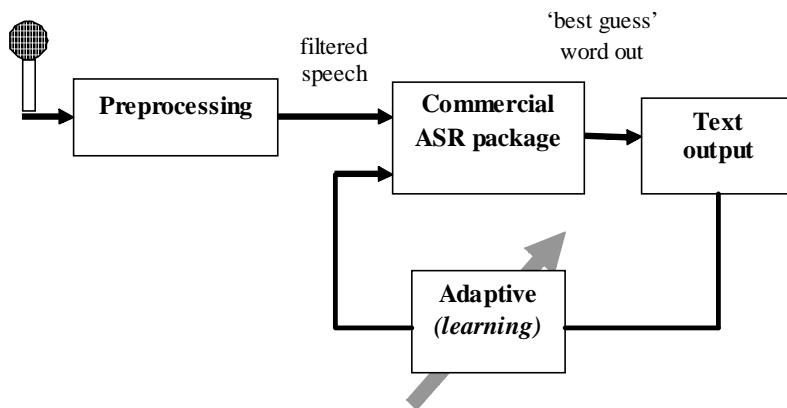
allows us to cease justifying our position through what we “know”. We are freed from defending our “positions” which in turn allows us to flow and imagine and become. It enables us to be freed from the shackles of what could be referred to as the death of a new beginning, since the discordant, rather than harmony, is valued.

The safety of the DTW process in the office “theatre” has allowed us to surrender and explore what could become, whether we think we will like it or not. The process allows us to become willing to let the unknown, the unacknowledged, the half thought, the unsubstantiated, the hunch to bubble up into awareness, be shared with and acknowledged by a trusted other within our community, and then if we choose take it to a public space. The theatre of the office provides a relationship space that is the equivalent of a well-articulated space, a space that forms as we bring ourselves into relationship with the known, the unknown, and the “not yet”, so that we are able to conceive and generate knowing and do what we would not otherwise have been able to do.

INAPPROPRIATENESS OF USING COMMERCIAL (*OFF-THE-SHELF*) SPEECH RECOGNITION SOFTWARE FOR DRAW-TALK-WRITE

Current best practice in assisting orally-proficient, literacy-inefficient speakers to become written-text literate is labour-intensive, and involves a learning facilitator working one-on-one with the speaker, as previously described. An *automatic* word recognition system is currently under development to enable people to impart their stories *directly* in text form, without the assistance of a human facilitator, in other words, automation of the DTW process. It should be emphasized that this system differs from virtually all other speech recognition systems in that users are not required to interact

Figure 14. Conventional ASR system



via *written* text (such as Kohonen, 1988); at the heart of the system is a neural network-based pattern recognizer which translates speech patterns into *visual* rather than textual cues.

Commercial speech recognition packages such as IBM's ViaVoice® for PC or Macintosh (www-306.ibm.com/software/voice/viavoice) and Dragon's NaturallySpeaking™ (www.dragontalk.com/) for MS-Windows™ typically adapt themselves to individual users during an initial training or “settling in” period (although some would argue that users need to adapt themselves to the idiosyncrasies of the software — sic). Such adaptation is predicated, however, on users being text literate, in other words, being able to first read, then respond to feedback presented to them by way of the computer's screen. Obviously such an approach is not viable for text-*illiterate* users, who nevertheless are typically *orally* proficient, and often come from a culture rich in imagery and story-telling. More appropriate feedback for such users would therefore be images rather than written text.

Automation of the DTW Process

The aim of this study is to replace the one-on-one human facilitator in the feedback loop (Figure 14 — adaptive learning) with an automated system — one based on a neural network pattern classifier. Further, our basic aim is no longer to produce text output from speech input, but rather images. Cast in terms of a pattern recognition problem, suddenly artificial neural networks become relevant again, whereas for conventional speech recognition they have been largely superseded by hidden Markov models.

A 2002 pilot study centered around the Macintosh-based proof-of-concept system shown in Figure 15. It comprised: (a) a voice input pre-processor (microphone, sound card, and noise filter), (b) a fast Fourier transform package (which converted sampled words to frequencies), and (c) an ANN pattern classifier (the output from which was the 1-of-*n* “best match” from the reference word look-up table). We hasten to add that this reference vocabulary was kept very small in this first instance.

Figure 15. Original Macintosh-based system (2002)

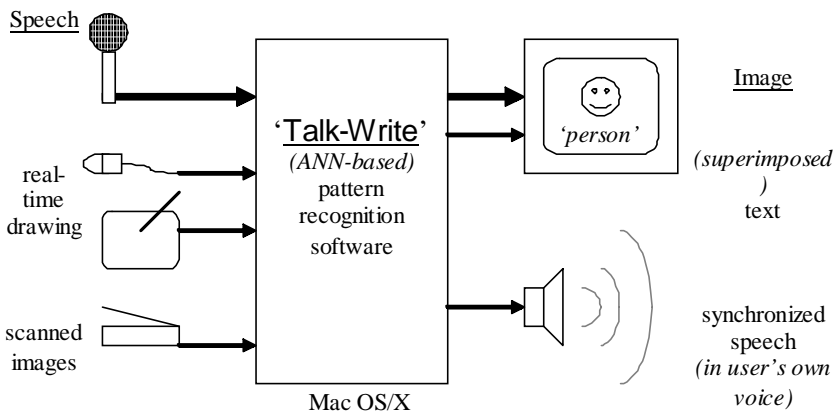
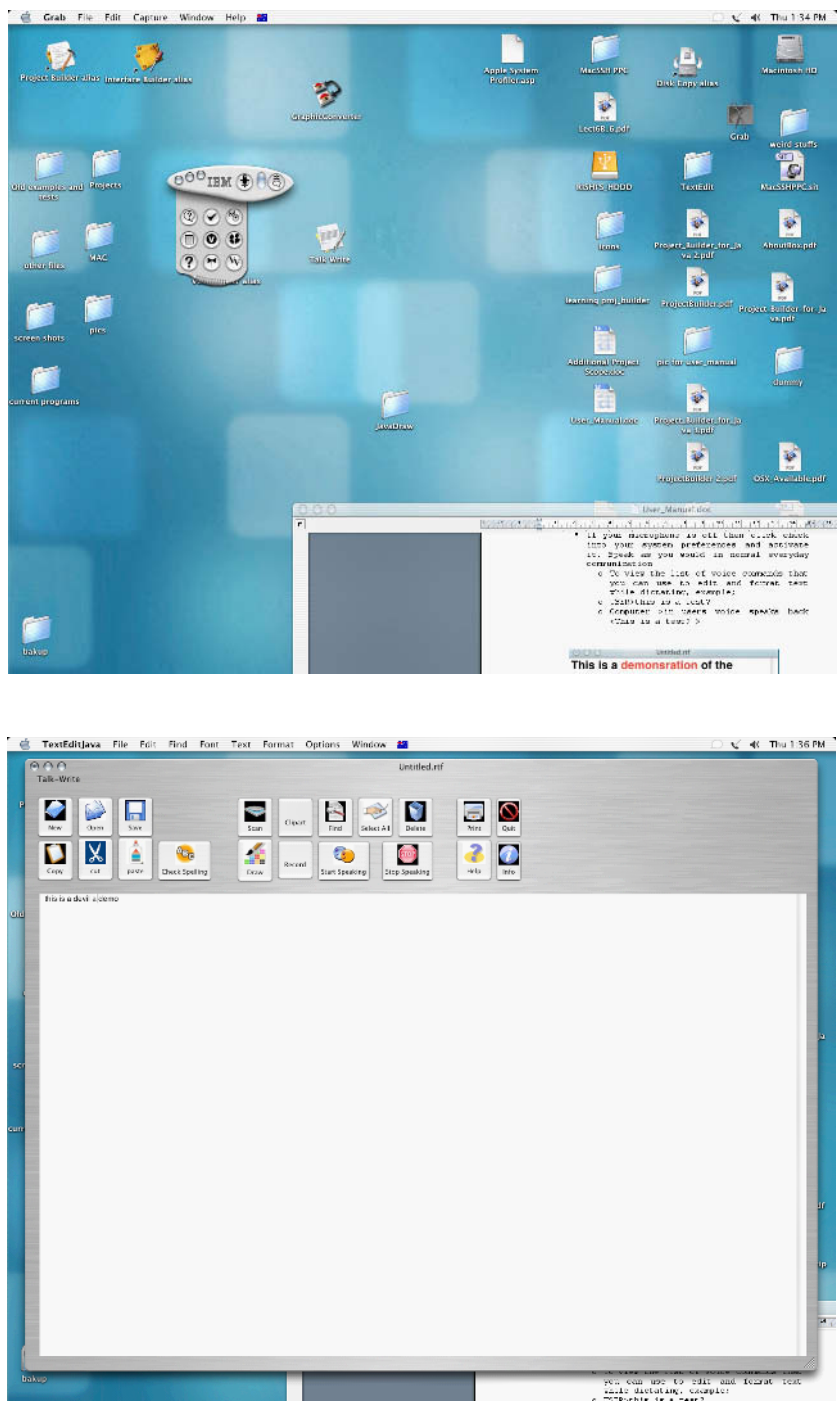


Figure 16. Apple Macintosh G4™ screen dumps of Talk-Write software (top: user manual; bottom: input)



By the end of this 12-month inaugural study, whilst some success was forthcoming with each of these three sub-sections, the overall system performance was somewhat lacking.

A second system was developed the following year. A screen dump from the Apple Mac G4 screen is shown in Figure 16, from which we see integration of IBM ViaVoice®, as well as support for additional input devices — namely scanner, graphics tablet and mouse. These latter devices are needed in order to augment speech input. More specifically, users are able to input their own drawings (either pre-prepared or new, via the tablet or mouse), in order to complement their oral stories.

As a first approximation to speech recognition for literacy, images could simply be linked on a one-to-one basis with words in the inbuilt vocabulary look-up table — whether that be as part of the Macintosh OS/X™ inbuilt speech library, or third-party software packages such as Dragon NaturallySpeaking™ or IBM ViaVoice® (the latter is shown in Figure 16). Ultimately however, we are aiming to do this the other way around — in other words, to produce image output from speech input, then link the former on a one-to-one basis to text. Over time the user begins to associate (internalize) these words and images as part of the DTW process.

Other system features critical to producing an automated DTW “engine” are:

1. storage of speech input in a form easily indexed and retrieved as needed, and
2. synchronised playback of keywords/phrases in the speaker’s *own* voice rather than in the unrealistic styles used in commercial speech synthesis packages.

Up to the present time, an unrealistically small reference vocabulary has been used; obviously this would need to be expanded significantly before a production version is released into the marketplace. More to the point, we have yet to determine just what constitutes a “minimum yet sufficient”-sized vocabulary to enable users to tell their stories (and no doubt this will vary considerably from user to user).

CONCLUSION

This work-in-progress has thrown up numerous exciting possibilities for future investigation. Apart from the system issues outlined above, there is much experimentation that could be performed to determine optimum pattern recognition configurations (to date, only simple, naïve multi-layer perceptron/back-propagation neural networks have been used). Likewise, we have yet to benchmark ANNs against alternative pattern classifier approaches.

The future possibilities and applications of draw-talk-write are limited only by our fears and lack of perceived safety. For example, “literary dramaturgy” has recently enabled people to consider and experiment public writing processes with literacy-inefficient people. DTW provides rich potential for minorities to voice, witness, and be heard by audiences who demand text and belittle those that have not mastered it.

What we need to assist us in our endeavours is technology that can record voice into text, synchronise it with playback in the voice of the narrator and the production of images, in other words, an intelligent system which incorporates word recognition, but which is configured in a manner that enables computer illiterate people to utilise the

system. Thus the computer system needs to respond to the user, rather than constrain people because they cannot meet the demands or limitations of the machine.

Lastly, successful automation of DTW on a computer platform would have far-reaching consequences beyond the specific (text-illiterate) section of the population of interest in the present study. Indeed, *any* community possessing a strong oral (story-telling) tradition could stand to benefit from this technology. Moreover, since the system output is images rather than text, it would have universal appeal.

ACKNOWLEDGMENTS

Financial support for this project from the Apple University Development Fund — AUDF — is gratefully acknowledged. The contributions of the following colleagues are likewise much appreciated: Kim Draisma, Ernie Blackmore, Marion Worthy, David Welsh, Frances Laneyrie, “the Dancer” and “the Prisoner” (fellow DTW travellers); Brian Pinch, Sunil Hargunani, Rishit Lalseta, Benjamin Nicholson, Riaz Hasan, Brij Rathore, and Phillip Dawson (software developers); and Professor Michael Wagner (proofreading and comments).

REFERENCES

- Ainsworth, W. A. (1997). Some approaches to automatic speech recognition. In W. J. Hardcastle & J. Laver (Eds.), *The handbook of phonetic sciences* (Chap. 24). Oxford, UK: Blackwell.
- Atal, B. S., & Hanauer, S. L. (1971). Speech analysis and synthesis by linear prediction of the speech wave. *Journal of the Acoustic Society of America*, 50(2), 637-655.
- Bourlard, H., & Morgan, M. (1994). *Connectionist speech recognition*. Boston: Kluwer Academic Press.
- Bourlard, H., & Wellekens, C. (1990). Links between Markov models and multilayer perceptrons. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(12), 1167-1178.
- Brunet, P. T., Pandya, A. S., & Pinera, C. V. (1994, June 27-July 2). Artificial neural networks for phoneme recognition. *Proceedings of the IEEE World Congress on Computational Intelligence*, Orlando, FL (pp. 4473-4478).
- Cambourne, B., & Turbill, J. (1987). *Coping with chaos*. Rozelle, NSW: Primary English Teaching Association.
- Campbell, J. P. (1997). Speaker recognition: A tutorial. *Proceedings of the IEEE*, 85(9), 1437-1462.
- Cosatto, E., Ostermann, J., Graf, H. P., & Schroeter, J. (2003). Lifelike talking faces for interactive services. *Proceedings of the IEEE*, 91(9), 1406-1429.
- Deng, L., & Huang, X. (2004). Challenges in adopting speech recognition. *Communications of the ACM*, 47(1), 69-75.
- Deng, L., & O'Shaughnessy, D. (2003). *Speech processing — A dynamic and optimisation-oriented approach*. New York: Marcel Dekker.
- Draisma, K., Gluck, R., Hancock, J., Kanitz, R., Price, W., Knell, J., Sharman, G., & Squires, J. (1994, October). Look what you get when the system is a bit more flexible and

- generous. *Journal of the Australian & New Zealand Student Services Association News*, 4.
- Ferguson, J. (1980). *Hidden Markov models for speech*. Princeton, NJ: IDA.
- Fritsch, J., Hild, H., Meier, U., & Waibel, A. (2000). Time-delay neural networks for NN/HMM hybrids: A family of connectionist continuous-speech recognition systems. In S. Katagiri (Ed.), *Handbook of neural networks for speech processing* (Chap. 7). Norwood, MA: Artech House.
- Fulcher, J. (1997). Image processing. In E. Fiesler & R. Beale (Eds.), *Handbook of neural computation* (Chap. F1.6). Oxford, UK: Oxford University Press.
- Fulcher, J. (2005). Secure, remote, ubiquitous access of electronic health records. In D. Bangert & R. Doktor (Eds.), *Human and organizational dynamics in e-health: A global perspective*. Oxford, UK: Radcliffe Medical Publishers.
- Fulcher, J., Gluck, R., Worthy, M., Draisma, K., & Vialle, W. (2003). A neural network, speech-based approach to literacy. *Proceedings of the International Symposium on Information & Communications Technology*, Dublin, Ireland, September 27-29 (pp. 27-32).
- Gallimore, R., & Tharp, R. (1990). Teaching mind in society: Teaching schooling and literate discourse. In L. C. Moll (Ed.), *Vygotsky and education*. Cambridge, UK: Cambridge University Press.
- Ganapathiraju, A., & Picone, J. (2000). Hybrid SVM/HMM architecture for speech recognition. *Proceedings of the 6th International Conference on Spoken Language Processing*, Beijing, China, October 16-20 (pp. 504-507).
- Gardner, H. (1983). *Frames of mind: The theory of multiple intelligences*. New York: Basic Books.
- Gauvain, J. -L., & Lamel, L. (2000). Large-vocabulary continuous speech recognition: Advances and applications. *Proceedings of the IEEE*, 88(8), 1181-1200.
- Giangola, J. P., Balogh, J., & Cohen, M. H. (2004). *Voice user interface design*. Boston: Addison Wesley.
- Gluck, R., & Draisma, K. (1997). The Pipeline Project: An Investigation of Local Infrastructures to Educate Aboriginal Professionals for Community, Professional and Industrial Organisations in the Illawara and Surrounding Regions. Evaluations and Investigations Program, Higher Education Division, Department of Employment Education, Training and Youth Affairs, Canberra, Australia.
- Gluck, R., Draisma, K., Fulcher, J., & Worthy, M. (2004). Draw-talk-write. *Proceedings of the 6th Biannual Language and Academic Skills in Higher Education Conference*, Flinders University, Adelaide, SA (pp. 109-116).
- Gluck, R., Vialle, W., Lysaght, P., & Larkin, J. (1998). What have I learned? In R. Corderoy (Ed.), *Flexibility: The next wave? Proceedings of the 15th Annual Conference of the Australasian Society for Computers in Learning in Tertiary Education - ASCILITE*, Wollongong, NSW (pp. 297-308).
- Gluck, R., Vialle, W., & Lysaght, P. (1999a). Reorienting mainstream learning through aboriginal experience. *Proceedings of The Challenge of Change in Education Conference — Change and Education Research Group (CERG)*, University of Technology, Sydney, February.
- Gluck, R., Vialle, W., & Lysaght, P. (1999b). Are you doing time or getting an education? *Proceedings of the 34th World Congress of the International Institute of Sociology: Multiple Modernities in an Era of Globalization*, Tel Aviv, July 11-15.

- Haton, J. -P. (2000). Probability-oriented neural networks and hybrid connectionist/stochastic networks. In S. Katagiri (Ed.), *Handbook of neural networks for speech processing* (Chap. 8). Norwood, MA: Artech House.
- Hornik, K. (1991). Approximation capability of multilayer feedforward networks. *Neural Networks*, 4, 2151-2157.
- Huang, X. -D., Acero, A., & Hon, H. -W. (2001). *Spoken language processing — A guide to theory, algorithms, and system development*. Upper Saddle River, NY: Prentice Hall.
- Huang, X. -D., Akiri, Y., & Jack, M. (1990). *Hidden Markov models for speech recognition*. Edinburgh, UK: Edinburgh University Press.
- Itakura, F., & Saito, S. (1970). A statistical method for estimation of speech spectral density and formant frequencies. *Electronics & Communications in Japan*, 53-A, 36-43.
- Jordan, M. I., & Jacobs, R. A. (1992). Hierarchies of adaptive experts. *Advances in Neural Information Processing*, 4, 985-993.
- Jordan, M. I., & Jacobs, R. A. (1994). Hierarchical mixtures of experts and the EM algorithm. *Neural Computation*, 6(2), 181-214.
- Juang, B. H., & Furui, S. (2000). Automatic recognition and understanding of spoken language — ZA first step toward natural human-machine communication. *Proceedings of the IEEE*, 88(8), 1142-1165.
- Katagiri, S. (2000). Speech recognition. In S. Katagiri (Ed.), *Handbook of neural networks for speech processing* (Chap. 3). Norwood, MA: Artech House.
- Keller, E. (1994). *Fundamentals of speech synthesis and speech recognition: Basic concepts, state-of-the-art and future challenges*. Chichester, England: John Wiley.
- Kohonen, T. (1997). *Self-organizing maps* (2nd ed.). Berlin: Springer-Verlag.
- Kotelly, B. (2003). *The art and business of speech recognition*. Berkeley, CA: Peachpit Press.
- Lang, K., & Hinton, G. (1988). *The development of the time-delay neural network architecture for speech recognition* (Tech. Rep. No. CMU-CS-88-152). Pittsburgh, PA: Carnegie-Mellon University.
- Lang, K., Waibel, A., & Hinton, G. (1990). A time-delay neural network architecture for isolated word recognition. *Neural Networks*, 3(1), 23-43.
- Leshno, M., Lin, V., & Schoken, A. (1993). Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural Networks*, 6, 861-867.
- Lippmann, R. (1989). Review of neural networks for speech recognition. *Neural Computation*, 1(1), 1-38.
- Lysaght, P. (2001). *Intelligent profiles: A model for change in women's lives*. Doctor of education thesis, University of Wollongong, Australia.
- Man, K. F., Tang, K. S., & Kwong, S. (1999). *Genetic algorithms: Concepts and designs*. London: Springer Verlag (Chap. 8 — GA's in Speech Recognition Systems).
- Marsic, I., Medl, A., & Flanagan, J. (2000). Natural communication with information systems. *Proceedings of the IEEE*, 88(8), 1354-1366.
- Masaki, S. (2000). The speech signal and its production model. In S. Katagiri (Ed.), *Handbook of neural networks for speech processing* (Chap. 2). Norwood, MA: Artech House.

- McDermott, E. (2000). Discriminative prototype-based methods for speech recognition. In S. Katagiri (Ed.), *Handbook of neural networks for speech processing* (Chap. 5.6 LVQ). Norwood, MA: Artech House.
- McDermott, E., & Katagiri, S. (1991). LVQ-based shift-tolerant phoneme recognition. *IEEE Transactions on Speech Processing*, 39(6), 1398-1411.
- Morgan, D. P., & Scofield, C. L. (1991). *Neural networks and speech processing*. Amsterdam: Kluwer.
- Morgan, N., & Bourlard, H. A. (1995). Neural networks for statistical recognition of continuous speech. *Proceedings of the IEEE*, 83(5), 742-770.
- Ney, H. (1984). The use of a one-stage dynamic programming algorithm for connected word recognition. *IEEE Transactions on Acoustics, Speech, & Signal Processing*, ASSP-32(2), 263-271.
- O'Shaughnessy, D. (2003). Interacting with computers by voice: Automatic speech recognition and synthesis. *Proceedings of the IEEE*, 91(9), 1272-1305.
- Oviatt, S. (2002). Breaking the robustness barrier: Recent progress on the design of robust multimodal systems. In M. Zelkowitz (Ed.), *Advances in computers* (pp. 305-341). San Diego, CA: Academic Press.
- Potamianos, G., Neti, C., Gravier, G., Garg, A., & Senior, A. W. (2003). Recent advances in the automatic recognition of audiovisual speech. *Proceedings of the IEEE*, 91(9), 1306-1326.
- Rabiner, L., & Juang, B. -H. (1993). *Fundamentals of speech recognition*. Englewood Cliffs, NJ: Prentice Hall.
- Ramachandran, R. P., & Mammone, R. (Eds.). (1995). *Modern methods of speech processing*. Norwell, MA: Kluwer Academic Publishers.
- Rieber, R. W., & Carton, A. S. (Eds.). (1987). *The collected works of L. S. Vygotsky*. New York: Plenum.
- Robinson, T., Hochberg, M., & Renals, S. (1996). The use of recurrent neural networks in continuous speech recognition. In C. -H. Lee, F. K. Soong, & K. K. Paliwal (Eds.), *Automatic speech and speaker recognition: Advanced topics* (pp. 233-258). Boston: Kluwer Academic Publishers.
- Rumi (c1250) A great wagon. Translated by C. Barks, J. Moyne, A. Arberry, & R. Nicholson (1995). *The essential Rumi*. New York: Harper Collins.
- Sakoe, H., & Chiba, S. (1978). A dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, & Signal Processing*, ASSP-26(1), 43-49.
- Schuster, M. (2000). Recurrent neural networks for speech recognition. In S. Katagiri (Ed.), *Handbook of neural networks for speech processing* (Chap. 6). Norwood, MA: Artech House.
- Sridharan, S., Leis, J., & Paliwal, K. K. (2000). Speech coding. In S. Katagiri (Ed.), *Handbook of neural networks for speech processing* (Chap. 4). Norwood, MA: Artech House.
- Stadermann, J., & Righoll, G. (2003). Flexible feature extraction and HMM design for a hybrid speech recognition system in noisy environments. *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing* (Vol. 3), April 6-10 (pp. 332-335).

- Thompson, R. (2003). *Stepping out with gargantua — Learning new research practices in the educational theatre of absurdity*. Doctor of education thesis, University of Canberra, Australia.
- Trentin, E., & Gori, M. (2003). Robust combination of neural networks and hidden Markov models for speech recognition. *IEEE Transactions on Neural Networks*, 14(6), 1519-1531.
- Vaughan-Nichols, S. -J. (2004). Voice authentication speaks to the marketplace. *IEEE Computer*, 37(3), 13-15.
- Van der Veer, R., & Valsiner, J. (1994). *The Vygotsky reader*. Oxford, UK: Blackwell.
- Vygotsky, L. S. (1934/1987). Thinking and speech. In R. W. Rieber & A. S. Carton (Eds.), *The collected works of L.S. Vygotsky, Volume 1: Problems of general psychology*. (Translated by N. Minick). New York: Plenum.
- Vygotsky, L. S. (1962). *Thought and language*. Cambridge, MA: MIT Press.
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. In M. Cole, V. J. Steiner, S. Scrivener, & E. Suberman (Eds.). Cambridge, MA: Harvard University Press.
- Waibel, A., Hanazawa, T., Hinton, G., Shikano, K., & Lang, K. (1989). Phoneme recognition using time-delay neural networks. *IEEE Transactions on Acoustics, Speech, & Signal Processing*, ASSP-37(3), 328-339.
- Waibel, A., & Lee, K. -F. (1990). *Readings in speech recognition*. San Mateo, CA: Morgan Kaufmann.
- Walker, W., Lamere, P., Kwok, P., Raj, P., Singh, R., Gouvea, E., Wolf, P., & Woepel, J. (2004). *Sphinx-4: A flexible open source framework for speech recognition* (Tech. Rep. No. TR-2004-139). SUN Microsystems.
- Wang, J. -F., Furui, S., & Juang, B. -H. (Eds.). (2004). *Real world speech processing*. Boston: Kluwer Academic Publishers.
- Wells, G. (1999). *Dialogical inquiry: Towards a sociocultural practice and theory of education*. New York: Cambridge University Press.
- Williams, R. (1997). *Problems in materialism and culture: Selected essays*. London: Verso.
- Wu, J., & Chan, C. (1993). Isolated word recognition by neural network models with cross-correlation coefficients for speech dynamics. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 15(11), 1174-1186.
- Young, S. J., Wood, P. C., & Byrne, W. J. (1996). Spontaneous speech recognition for the credit card corpus using the HTK toolkit. *IEEE Transactions on Speech & Audio Processing*, 2(4), 615-621.
- Zhang, M., Xu, S., & Fulcher, J. (2002). Neuron-adaptive higher order neural network models for automated financial data modelling. *IEEE Transactions on Neural Networks*, 13(1), 188-204.
- Zwicker, E. (1961). Subdivision of the auditory frequency range into critical bands. *J. Acoustic Society of America*, 33, 248.

ENDNOTES

- ¹ Two well known examples of ANNs learning unexpected input-output associations are (a) sunny versus cloudy days instead of images of forests with and without Army tanks, and (b) photographs of ‘males’ versus ‘females’ being classified not on the basis of gender, but rather on the amount of white space between the tops of their heads and the top of the photograph [ref. The Dream Machine, Episode#4, BBC 1991].
- ² Technical And Further Education — technical/vocational post-secondary school colleges.

APPENDIX: AUTOMATIC SPEECH RECOGNITION RESOURCES

I — Speech Research:

- <http://research.microsoft.com/srg/> (*Redmond, USA*);
<http://research.microsoft.com/speech/> (*China*)
- www.almaden.ibm.com/
- www.speech.cs.cmu.edu (*Carnegie Mellon University*)

II — Speech Resources:

- <http://cmusphinx.sourceforge.net/html/cmusphinx.php> (*Linux, Unix & Windows*)
- <http://online.ldc.upenn.edu/> (*Linguistic Data Consortium, University of Pennsylvania*)
- <http://emu.sourceforge.net/release.shtml> (*Emu = Tcl/Tk-based Linux/Solaris, + windows/Mac binaries*)

III — Speech Databases:

- www.nist.gov/speech/tests/spk/2000/ (US Department of Commerce) — see also The DARPA TIMIT Acoustic-Phonetic continuous speech corpus (CD-ROM) NIST SpeechDisk 1-1.1, NTIS order# PB91-505065, 1990
- www.speech.cs.cmu.edu/comp.speech/
- <http://isw3.aist-nara.ac.jp/IS/Shikano-lab/database/internet-resource/e-www-site.html#Speech%20Database>
- <http://mambo.ucsc.edu/psl/speech.html>

IV — Other Speech-Related Site of Interest:

- www.biometrics.org/html/research.html (*speech & other biometrics*)

Chapter V

Smart Cars: The Next Frontier

Lars Petersson, National ICT Australia, Australia

Luke Fletcher, Australian National University, Australia

Nick Barnes, National ICT Australia, Australia

Alexander Zelinsky, CSIRO ICT Centre, Australia

ABSTRACT

This chapter gives an overview of driver assistance systems (DAS) in general and the Smart Cars project in particular. In the Driver Assistance Systems Section, a set of key competencies for an effective DAS are identified by comparing with a human co-pilot, namely, traffic situation monitoring, driver's state monitoring, vehicle state monitoring, communication with the driver, vehicle control, and a reasoning system. It is also recognised that such a system must be intuitive, non-intrusive and override-able. A few of the currently available commercial systems are mentioned in the following section. The Smart Cars project, which is a joint project between the Australian National University and National ICT Australia, is then introduced. A number of different research directions within the project are then presented in detail: obstacle detection and tracking, speed sign detection and recognition, pedestrian detection, and blind spot monitoring.

Figure 1. The research platform, a 1999 Toyota Land Cruiser™, equipped with the appropriate actuators, sensors, and computing power to perform monitoring and control of the vehicle



INTRODUCTION

The number of cars in use on our roads increases every year, and with that, the number of accidents. However, by introducing systems and technologies that help the driver in difficult or dangerous situations, the car industry has been able to keep the number of fatal accidents down. Examples of such systems are ABS-brakes and smart air bag deployment. New, more advanced technologies that can help the driver have, in recent years, started to be deployed in production vehicles. Examples of such systems are parking aids, lane departure warning systems, and emergency brake systems. The research in the Smart Cars project is particularly concerned with such advanced driver assistance systems, namely, that assist the driver in controlling the car, but keep the driver in the loop. Impressive work in this and related areas has been performed by Dickmanns and Zapp (1987), Dickmanns (1999, 2000), Broggi, Bertozzi, and Fascioli (2001), Mertz, McNeil, and Thorpe (2000), Zhao and Thorpe (2000), Aufere, Gowdy, Mertz, Thorpe, Wang, and Yata (2003), and Bertozzi, Broggi, Carletti, Fascioli, Graf, Grisleri, and Meinecke (2003). Their work deals to a large extent with the sensing aspect of driver assistance, which is essential to create robust and reliable systems. An interesting research area is, however, how to handle the information flow generated. Depending on the context, information has different significance. For example, how are warnings most efficiently conveyed?

DRIVER ASSISTANCE SYSTEMS

A driver assistance system (DAS) may perform activities like relieving the driver of distracting routine activities, warn about upcoming situations, and possibly take control of the car if an accident is imminent. Depending on the task to be performed, a DAS must have appropriate levels of competencies in a number of areas. If we consider the DAS from the perspective of a human co-pilot, it is easier to pick out the important aspects.

To be of assistance, the co-pilot needs to be aware of what is going on outside of the car; for example, are there any pedestrians in sight, where are they going, how is the road turning, and so on. Moreover, we would like our co-pilot to warn us if we have not noticed an upcoming situation. That means that not only should the co-pilot be aware of what is going on outside of the car, but also what is happening inside, in other words, the driver's responses. In addition, our co-pilot must know where the vehicle is going, how fast it is going, if we are braking, accelerating, and so forth, in order to make good decisions. Good decisions are a result of good reasoning. A successful driver/co-pilot team requires good communication. The co-pilot must not be intrusive or present the driver with too much information. Finally, if the co-pilot notices that the driver does not respond to a situation that will result in an accident, they must be able to take control of the car.

Returning to our non-human (automated DAS) co-pilot, we can condense the above to the following key competencies:

- Traffic situation monitoring
- Driver's state monitoring
- Vehicle state monitoring
- Communication with the driver
- Vehicle control
- Reasoning system

The first three collect information which the DAS can use to analyse the current situation. The fourth — communication with the driver — provides both input to the DAS and output to the driver. For example, the driver can specify an overall goal, or the DAS can give information to the driver. Vehicle control is necessary if it is expected that the DAS should be able to perform any semi- or fully-autonomous manoeuvres. A reasoning system may range from a direct mapping from an input to an output, to a complex system using the latest advances in artificial intelligence. The level of competence in each category is dependent on the specific task to be solved.

Finally, with a human co-pilot, the DAS should possess the following behavioural characteristics:

- **Intuitive** — The behaviour of the DAS must make immediate sense in the context of the standard driving task.
- **Non-intrusive** — It must not distract or disrupt the driver unless it is necessary.
- **Override-able** — The driver has ultimate control and can refuse assistance.

COMMERCIAL SYSTEMS

The car industry and related companies are quickly moving towards more complex systems to deploy in production vehicles. These range from non-critical systems such as automatic parking to safety critical systems like emergency brakes.

Toyota recently released an automatic parallel parking aid that comes as an option on their Prius™ model. The system uses a rear camera that views the potential parking spot, and the driver selects an appropriate area within that view by moving lines in the

image. After confirming the intended spot, the driver keeps their foot lightly on the braking pedal, and the car will start backing up, steering automatically. The system is still in its early stages of development, and can, so far, only park when it is possible to continuously back into a parking space. Although it currently has obvious limitations, it clearly shows that major car manufacturers are willing to deploy highly complex aids in order to remain competitive.

In the area of active safety, Honda has introduced a Collision Mitigation Brake System (CMS), in its Inspire™ model. The system predicts rear-end collisions and assists brake operation to reduce the impact to occupants and the vehicle itself. A millimetre-wave radar is used to estimate distances to vehicles ahead, relative speeds, and expected paths. In the event of a likely collision, there is an audible alarm, tightening of the seatbelt, and a brake assist function that compensates for insufficient pedal pressure to reduce the speed of impact. A system like this is clearly a safety critical system which can cause accidents by itself if triggered by mistake.

There are also commercial systems that monitor the actions of the driver like faceLAB™ from Seeing Machines (<http://www.seeingmachines.com>). faceLAB™ can track the pose of the head and the direction of the eye gaze, and can measure the amount of time the eyelids are closed. Volvo and other car manufacturers plan to incorporate this into their advanced driver assistance systems.

THE SMART CARS PROJECT

The Smart Cars project is a joint project between the Australian National University and National ICT Australia. The project was initiated in 2000 by Professor Alexander Zelinsky within the Research School of Information Sciences and Engineering. A research platform was built using a 4WD Toyota Landcruiser by equipping it with sensors and computing hardware, and by modifying the steering, braking, and accelerator. In our current research, we are particularly concerned with advanced driver assistance, systems that assist the driver in controlling the car, but keep the driver in the loop. This paradigm addresses, to some extent, the robustness issue. Completely autonomous systems need to be extremely robust and reliable — errors can be fatal — whereas an assistive system enhances the driver's existing capabilities. Using both external and internal sensing, a natural step is to fuse the available data and suggest a suitable response.

The focus of this project is to further identify effective methods for advanced driver assistance, to develop particular sensing, detection, and human-machine interface systems, and to make them robust and reliable.

Obstacle Detection and Tracking

Reliable obstacle detection and tracking has proved a challenging problem in automotive research. A moving sensor, varied lighting, and unknown object appearance preclude many classic segmentation techniques such as background subtraction or appearance-based feature detection. In response, progress has tended along two avenues: superior sensing and constraining the problem. Superior sensors have included laser range finders, millimetre wave radars, or large baseline stereo cameras


(Broggi et al., 2001; Grover, Brooker, & Durrant-White, 2001; Roberts & Corke, 2000). Constrained problem-based solutions have used significant assumptions like flat road models, featureless road surfaces, or looking for known/car-like objects (Dellaert, Pomerleau, & Thorpe, 2001). Our approach takes the middle ground with standard sensors and some weak constraints on the problem. The strength of the technique is that it combines several sensor data processing techniques and some reasonable assumptions about the obstacles (such as a consistent size and location over time) to develop a robust detection and tracking system. A key strength of the system is that additional information sources (like better sensors, or image processing algorithms) can easily be added, to improve the overall performance or to handle a particular case, without modification to the existing system.

System Overview

The obstacle detection and tracking system has three phases of operation: detection, distillation, and tracking.

Figure 2 shows these phases which operate concurrently, detecting new obstacles while tracking previously-detected obstacles. The first phase uses a set of “bottom up”, whole- image techniques (stereo disparity, optical flow, colour consistency) to search the image space for likely obstacles. The second phase uses the particle filter-based “distillation_algorithm” to provide hypothesis-based or “top-down” processing on the results of the first phase — the “distillation algorithm” has also been used for face tracking (Loy, Fletcher, Apostoloff, & Zelinsky, 2002) and lane tracking (Apostoloff & Zelinsky, 2003). Sets of particles representing each potential obstacle are injected into the filter state-space in a Gaussian distribution around the detected obstacle location. Particles representing unsubstantiated obstacles dissipate. The remaining potential

Figure 2. Phase 1. obstacle detection, phase 2. obstacle distillation, phase 3. obstacle tracking



obstacles are tracked within the distillation framework between frames. Found obstacles are represented by clusters of particles which remain consistent over time. The centroid and spread of these clusters are then used to bootstrap an extended Kalman filter to explicitly track each obstacle.

Obstacle Detection

The obstacle detection phase is based on the coarse segmentation of potential obstacles from a stereo template correlation-based disparity map and image gradient-based optical flow data. As mentioned in Franke and Heinrich (2002), the range of optical flow values and disparities encountered in the road scene is large. Image pyramids are used to compute fast disparity and flow maps across a large range of disparities. For the case of optical flow, we implement a method similar to Simocelli (1999). The optical flow is computed for the most coarse images, then the result is used to warp the next higher image resolution to maintain an acceptably small image motion at each level. The penalty for using a coarse-to-fine approach is that any errors occurring at any image resolution are propagated and amplified into the finer images.

For the disparity map estimation, image pyramids give a couple of added benefits in addition to increasing the range of disparities estimated; in the case of the disparity map estimation, no image warping between resolutions is performed. Correlation-based techniques are plagued with the issue of using the correct template window size. A large

Figure 3. Top: greyscale image; bottom: disparity map with road surface removed (dark is far, bright is close; the car on the left is across a disparity range from 15-23 pixels, the car on the right at seven pixels, cars in the distance are 4-6 pixels)





Figure 4. Road scene with obstacles identified with bounding boxes

window allows a more reliable match, but causes overly-smooth disparity maps. A small window size allows for finer features to be represented, but introduces noise due to erroneous matches. Fortunately, obstacle detection in road scenes usually supports the rule-of-thumb that close objects are large, and distant objects, such as vehicles further along the road, are small. Using an image pyramid and calculating the disparity for each image resolution with the same size correlation window means that the latter is effectively halved for each image resolution, going from coarse to fine. This property is what we would prefer to match: large objects at large disparities, and smaller objects at small disparities (like near the horizon). As no warping is done between image resolutions, we can avoid the propagation of errors between image resolutions. At higher resolutions we are interested in finding distant objects with small disparities, whereas larger objects such as close vehicles are recovered at coarser image resolutions. One issue that arises is that coarse resolution images can only resolve disparities to half the accuracy of the next higher image resolution. As this works in opposition to the property of disparity estimates deteriorating as a function of the inverse of the distance, the effect on the resultant disparity map is acceptable.

The stereo disparity map uses an iterative box filtering technique developed by Faugeras, Hotz, Mathieu, Viéville, Zhang, Fua, Théron, Moll, Berry, Vuillemin, Bertin, and Proy (1993) and implemented as described by Kagami, Okada, Inaba, and Inoue (2000).

The optical flow and stereo data is overlaid, and a series of simple heuristics are used to remove noise, outliers, and background objects. The stereo data is further processed by removing the road surface using the V-Disparity technique developed by Labayrade, Aubert, and Tarel (2002). The disparity map is accumulated per row into a 2D

histogram. The ground plane is assumed to be dominant, making a contour in the histogram. The contour is approximated by line segments, and the road surface can then be subtracted from the disparity image. Figure 3 shows a road scene with the segmented disparity image underneath. The obstacles are visible once the road surface has been removed. Potential obstacles are extracted from this image using basic constraints of a minimum height and maximum height. Figure 4 shows the resultant segmented road scene. Note that several false positives exist. This is acceptable as the desired result of this phase is a low number of false negatives (missed obstacles) at the expense of some false positives (phantom obstacles), which will be filtered in the second phase of the system.

Obstacle Distillation

In this phase the potential obstacles identified above are “distilled”, using the “distillation algorithm”, (Loy et al., 2002) into consistently-detected obstacles. The “distillation algorithm” (Figure 5) is a combination of a particle filter with an intelligent cue processing system. The cue processing system changes the rate at which different sensor data is incorporated into the particle filter based on how well the sensors are performing. For example, stereo data may be disrupted by a momentary occlusion of one camera, in which case the information from this cue is ignored in favour of other cues which are unaffected.

Sets of particles representing each obstacle candidate are injected into the state-space in a Gaussian distribution around the potential obstacle’s detected location. Stereo disparity, optical flow, and colour consistency cues are again used to evaluate the potential obstacles; this time, however, only the projected locations of the particles are evaluated, not the whole image. Over several iterations of the filter, particles representing unsubstantiated obstacles dissipate. The remaining particles clump into obstacles which are tracked within the particle filter between frames. Found obstacles are represented by clusters of particles which remain consistent over time. Each cluster of particles that survives a minimum number of iterations is then checked against a Gaussian distribution at its centroid. If the Gaussian distribution adequately describes the cluster or particles, an extended Kalman filter-based tracker is initialised (phase three). Figure 6 illustrates clusters of particles detected representing obstacles to be replaced with Kalman filters.

Figure 5. Distillation — a visual cue processing framework





Figure 6. *Obstacle distillation: Uni-modal clusters in the particle filter are extracted for tracking.*



Obstacle Tracking

In the third phase of the system, an extended Kalman Filter and image template correlation is used to track each obstacle independently. Using the obstacle location extracted from the previous phase, a uniqueness detector is used to identify good image correlation templates from each of the stereo cameras. This collection of templates is then used to track the obstacle. The correlation templates are tracked independently in the image using normalised cross correlation. The collection of templates associated with each obstacle is then evaluated by using the mean shift in the image and correlation value. Templates tracking inconsistently or unreliably are discarded. The remaining templates are used to estimate the new location of the obstacle fed to the extended Kalman filter. The extended Kalman filter tracks the location of the vehicle in the 3D road coordinate system and uses a constant velocity motion model. The size of the vehicle estimated in the previous phase is assumed to be constant. Eventually the obstacle tracked is lost: either overtaken, obscured, too far in front of the vehicle to be seen, or any other random failure. This condition is identified by either no reliable image templates remaining or by a divergence in the co-variance matrix of the filter. In either case, the system will discard the extended Kalman filter and, as a precaution, inject a cluster of particles at the final location of the object back into the particle filter in the above phase of the system.

Figure 7 shows the output of the obstacle tracking engine. Each car is tracked using an independent extended Kalman filter and image template correlation. Later in this image sequence, the centre car is lost due to a template tracking failure (only one template is tracking reliably at this stage), then the second phase of the system quickly detects the vehicle again and tracks it using a new filter and new image templates. Also in this sequence, the car on the far right is obscured by an overtaking vehicle. Again this obstacle is lost, and the overtaking vehicle is detected and tracked instead.

Figure 7. Output of obstacle tracker; rectangles indicate obstacle bounding boxes (“+” indicate correlation template locations; “” indicate centroid of obstacle)*



Speed Sign Detection and Recognition


The task we are looking at here is, with the help of camera(s), to automatically detect and recognise road signs. Automatic road sign detection and recognition is a valuable help to the driver since the system can, for example, make the driver aware of the current speed limit or an upcoming stop sign.

Solving this problem, it makes sense to use the fact that the roadway is well structured. Under Australian law, the appearance of road signs is highly restricted. They must be of a particular size, and be a white sign with black numbers surrounded by a red circle. Unless the sign has been tampered with, signs will appear approximately orthogonally to the road direction. Finally, signs are always placed to be easily visible, so the driver can see them without having to look away from the road.

The system must be able to work at frame rate (30Hz) which poses constraints on the methods to be used. A direct approach is to apply normalised cross correlation to the raw traffic scene image; however, this approach is computationally prohibitive, although there are methods to reduce this slightly (Betke & Makris, 1994).

Dealing with the problem of extensive computation, we would like to discard most parts of the incoming images, thereby only performing expensive calculations on relevant areas. One approach to do this is to divide the problem into two separate problems, namely detection and classification. By using a cheap detection stage, we are able to perform a more expensive classification stage. Many approaches have introduced separate stages for sign detection and classification of different types of signs, for example, Priese, Klieber, Lakmann, Rehrmann, and Schian, (1994); Miura, Kanda, and

Figure 8. A typical candidate sign detected by the fast radial symmetry detector (a) at the size it appears in the image, and (b) close-up. The outer circle and numbers are narrow. Despite its consistent appearance as a small image to our eyes, it contains few pixels that could be said to be red, black, or even white.



Shirai (2000); Paclik, Novovicova, Somol, and Pudil, (2000); and Johansson (2002). This is particularly useful when a large number of sign types are to be classified. We argue that this can be an effective means of managing computation for even a small number of sign types if a detection stage is available that has low computational cost, facilitating real-time operation. It allows computationally intensive classification to be performed on only a small part of the input image stream, without requiring assumptions about where signs may appear.

Detecting candidate regions by using colour segmentation is the most common method for the detection stage. However, since the different colours (red, black, and white) of the texture in the sign do not occupy more than a couple of pixels next to each other, the bleeding of colours between pixels makes it virtually impossible to find any consistent colours (see Figure 8). Dealing with this problem could easily remove the fast segmentation advantage of the two-step approach and requires further research.

Another approach to detection is *a priori* assumptions about image formation. At its simplest, one can assume that the road is approximately straight, so large portions of the image can be ignored as signs will not appear in them. However, such assumptions can break down on curved roads, or with bumps such as speed humps. A more sophisticated approach is to use some form of detection to facilitate scene understanding, and thus eliminate a large region of the image. For example, Piccoli, De Micheli, Parodi, and Campani (1996) suggest large uniform regions of the image that correspond to the road and sky, and thus only looking in the region alongside the road and below the sky where signs are likely to appear. However, this will not be adequate in more difficult road scenes, such as shown in Figure 9(b). They also suggest ignoring one side of the image as signs will only come up on one side. This is inadequate for multi-lane highways, and is throwing away information in scenes such as Figure 9(a).

We propose a new efficient method for sign detection: the fast radial symmetry detector (Loy & Zelinsky, 2003). It is applicable to signs with a circular feature, a significant sub-set of signs. Many shape detectors are non-robust because they require closed shapes. Robust techniques such as Hough circle detection (Minor & Sklansky, 1981) are slow to compute over-large images. The fast radial symmetry detector is efficient enough to be run as a detector at frame rate. We are able to eliminate the vast

Figure 9. Some sample images with speed signs present (The quality of sign and the lighting varied within our sequences, along with the scale of the sign that appeared. Also, more than one sign may appear in a single image)



majority of false positives by considering only radially symmetric regions that are stable across several images, and have a high count of pixels relative to the radius.

We then apply cross correlation to the small number of candidates. This will be over a restricted part of the image. For cross-correlation, scale is a problem, typically requiring multiple templates at different resolutions. However, from the radius returned from the fast radial symmetry detector, we know the approximate scale of the template.

The following presents preliminary work on a system for visual speed sign recognition. It identifies possible candidate speed signs in an image stream from a video camera mounted within the car looking along the road. It then classifies the sign as to what speed the sign indicated. This information can then be compared against the vehicle's speedometer, and passed onto the driver if it appears that they have not reacted to the change of conditions. Our approach exploits the structured nature of the road to facilitate fast processing. Currently, the system has only been evaluated on 40 and 60 Km/hr signs.

Figure 10. Some sample images showing how the fast radial symmetry algorithm is able to detect circles in images



Candidate Detection

The fast radial symmetry detector (Loy & Zelinsky, 2003) is a variant on the circular Hough transform that executes in order kp , where p is the number of pixels, and k is the number of discrete radii that are searched. This is in contrast to the traditional circular Hough transform that executes in order kbp . For the traditional circular Hough transform, each edge pixel votes on all circles over a discrete set of radii k that could pass through that edge pixel. The factor b comes from discretisation into a number of bins on the gradient of circular tangents that could pass through this point. The fast radial symmetry detector eliminates the factor b by taking the gradient of the edge point directly from the output of the Sobel edge detector. In this way, the computation of the radial symmetry

detector is reduced by a factor of b , but also the resulting circle map is simplified by one dimension. This makes it suitable for real time use, for example, 13.2 ms for a 240 x 320 image (Loy & Zelinsky, 2003). Figure 10 shows the algorithm at work in some sample images.

Consistency checking can be performed over time. A circle must appear for at least a certain number of concurrent frames; it cannot have changed radius too much, and it must be in the same region of the image. It may be possible to model car motion in predicting the new location of the image where the sign should appear. However, one would have to assume a smooth straight road, which is quite restrictive, or include feedback of steering direction, which is complex, but may contribute further. The output of our detection phase is the region of the image in which the candidate appears, a radius — representing the scale of the candidate — and a centre point of the candidate.

Classification

We apply standard normalised cross-correlation. A template was taken from a real image for each of the possible sign numbers. There is some variance in the appearance of the text on speed signs, so it would be desirable to take a number of candidates showing different text styles; currently, we just have a single basic template for each speed. The template taken consisted of simply the actual numbers. We selected a large clear sign for the template. This was then scaled down in size by linear interpolation across all colour channels to form a total of eight scales for each template. At run time, a sub-set of these templates, chosen on the basis of candidate radius, was compared against the candidate. The maximum value of these was selected as the figure representing the quality of the match to a template.

As the sign is assumed to be vertical and parallel to the image plane, no scaling or rotation of the template is necessary. Further, as the centroid of the candidate is known, the approximate location where the numbers would be if they were present is also known. We perform the cross-correlation with the centroid moving over a small region for each template. These efficiencies result in a fast classification stage, again facilitating real-time processing.

Results

The system was run over several raw image sequences taken from the experimental vehicle (Figure 1). The sequences come from cameras in a binocular head located approximately in the position of the rear-view mirror. All images used in the experiments were taken of signs on public roads around Canberra, including on the Australian National University campus. Some of the sequences were taken at frame rate while driving at around the speed limit, while others were taken while stationary in the normal position on the road in front of the signs.

The two phases were evaluated independently. The radial symmetry detector was run over a total of 1,107 frames from the camera. From this sequence, 152 sign candidates were detected. Of these, 90% were correctly detected. This is quite a good number from a fast early classification stage. For most scenes, as the sign approached, it was detected many times before it passed out of view, which would give the classification sub-system a good chance to correctly classify the sign. For the clear scenes, the sign was detected for most frames where it appeared. Our results indicate that requiring a classification to


be consistent for several frames would eliminate most of the false positives generated by the symmetry detector.

Each candidate that was returned was individually evaluated for being either a 40 or a 60 sign. From 126 valid candidates, 96% were correctly classified. Of these, unfortunately, all incorrect classifications were in a single class — 40 signs. Of the 25, 40 signs returned as candidates, 75% were correctly classified. This is still a highly promising result. Classification of a sequence of images could be improved by requiring temporal continuity of classification before accepting a candidate. This is quite acceptable within a driving situation as we would expect to be able to view a sign for at least several frames if we are processing at frame rate.

Pedestrian Detection

On our roads, traffic accidents involving pedestrian-vehicle collisions cause significant fatality and injury. Automotive manufacturers will soon be required to meet certain impact ratings for pedestrian-vehicle collisions (Fredriksson, Haland, & Yang, 2001). To achieve these requirements, manufacturers are considering the inclusion of active pedestrian protection systems (PPS) such as rising engine hoods, pedestrian protection airbags, and providing warning to the driver. Such active systems require knowledge of pedestrian presence for correct activation and deployment. We have developed a prototype to fulfil the sensory needs for PPS.

Figure 11. The distinction between critical and non-critical pedestrians (we aim to detect the critical pedestrians)



Considering the automotive application for this work, the sensory system should achieve more than just pedestrian detection. Sufficient pedestrian information should be provided to enable intelligent decisions regarding pedestrian presence for correct PPS deployment. Once detected, a pedestrian needs to be localised and tracked to provide the necessary pedestrian information:

- pedestrian location
- pedestrian size
- pedestrian velocity
- certainty of prediction

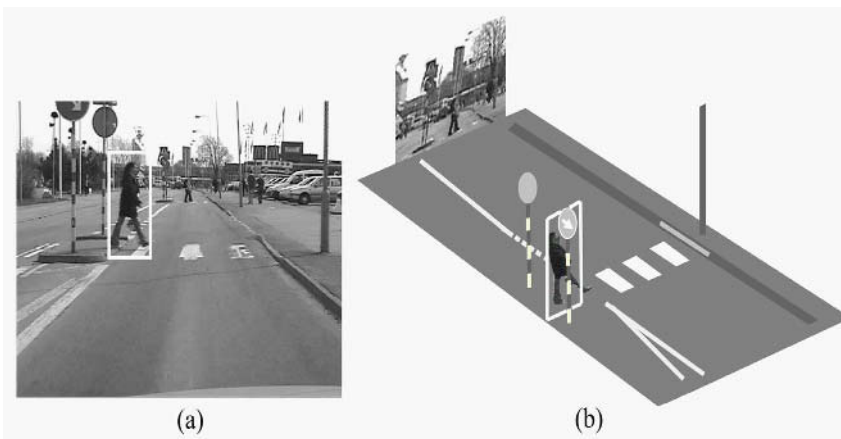
Furthermore, a crucial requirement regarding the automotive application is to achieve low false detection rates since the consequences of false detection are significant. To fulfil such needs, we use a vision-based, 3D temporal approach for detecting and tracking pedestrians.

Our system aims to detect at least the *critical* pedestrians in front of the vehicle; critical pedestrians are regarded as those who could be injured by the vehicle (see Figure 11).

Previous work in this area has utilised a range of sensing technologies including laser (Fuerstenberg & Lages, 2003), radar (Gavrila, 2000), and computer vision. Computer vision is showing the greatest potential, primarily because of its higher spatial resolution, enabling better discrimination between pedestrians and non-pedestrians. The approaches can be categorised as either 2D or 3D. Figure 12 depicts the fundamental difference between a 2D and 3D approach.

A 2D image analysis aims to recognise patterns in an image which resemble that of a pedestrian. These 2D approaches scan the entire image space, which is slow and, we

Figure 12. The difference between (a) 2D and (b) 3D approaches to the problem of pedestrian detection



believe, is prone to false detections. Furthermore, a 2D approach does not provide additional pedestrian information for use with PPS; robust range estimates cannot be provided and, therefore, such systems merely provide a detection/non-detection result. However, despite the drawbacks of such an approach, the development of 2D pattern analysis techniques to actually recognise pedestrians has made significant progress. These pattern analysis techniques include methods to detect pedestrian shape and walking motion.

We believe a 3D approach to pedestrian detection will achieve lower false detection rates since greater scene understanding is gained. Furthermore, a 3D approach inherently generates pedestrian information useful for PPS. Consider the simple example of a roadside billboard depicting a human. Most 2D methods would erroneously detect the human on the billboard as a pedestrian who could be injured by the vehicle. 3D information is necessary to realise the human image belongs to a larger structure.

System Architecture

As mentioned above, to reach the necessary low detection rates and real-time performance necessary for PPS, rather than attempting a “brute force” 2D image scan, we use a 3D temporal approach. Stereo vision is used for perceiving the environment in front of a host vehicle, thus providing 3D scene information. The camera setup is shown in Figure 13.

The software consists of three components: obstacle detection, obstacle classification, and pedestrian tracking (Figure 14).

Obstacle Detection

General obstacles are detected by segmenting a 3D representation of the scene in front of the host vehicle. Such 3D representation is obtained from a disparity map created from the stereo image pair. We generate dense disparity maps using the sum-of-absolute-differences (SAD) algorithm. Prior to disparity map generation, the image pair is rectified to ensure good correspondence matching. Objects are segmented from the disparity map using the ν -disparity algorithm (Labayrade, Aubert, & Tarel, 2002), which is a robust, fast, and accurate method for segmenting noisy disparity maps. The method provides scene understanding by recovering the ground surface and recognising which objects are sitting on this surface. Furthermore, ν -disparity is well suited to a moving camera platform since few assumptions regarding the ground surface in front of the host vehicle are made (in other words, ground plane calibration is not required). Therefore, segmentation is still achieved when the vehicle is approaching both inclined and declined road surfaces. Figure 15 illustrates the operation of the ν -disparity object detection algorithm.

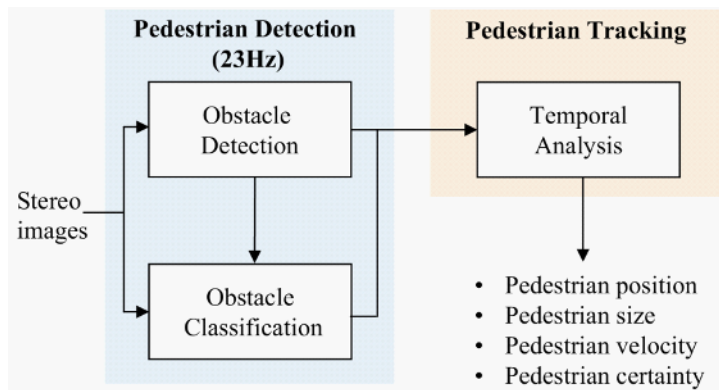
Obstacle Classification

Once segmented, each obstacle is classified as either pedestrian or non-pedestrian based on pedestrian shape using support vector machines (SVM). However, prior to such classification, we use a simple heuristic based on object size to eliminate objects not resembling the size of the human figure (for example, a vehicle is easily disregarded). Adult- and child-sized pedestrians are used to set the limits for object size passed by the heuristic, with a tolerance on the limits to ensure that pedestrians whose size is altered by accessories (for instance, an umbrella) are not rejected. This heuristic is not only a

Figure 13. System hardware - the host vehicle is equipped with a stereo camera configuration, a PC, and appropriate power supply accessories



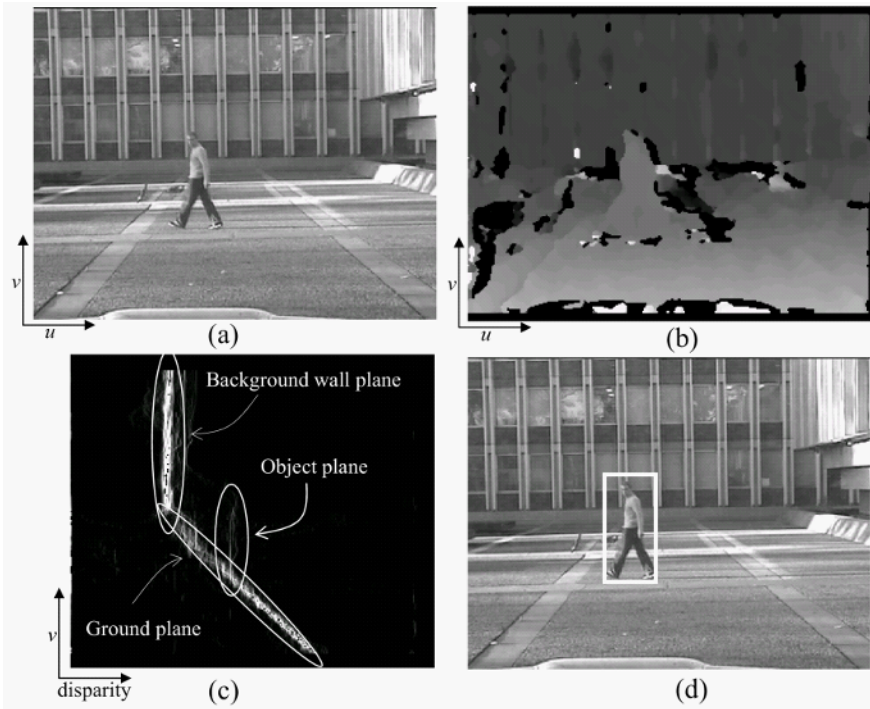
Figure 14. Pedestrian detection software structure



fast initial classification, but it eliminates the possibility of these objects being incorrectly classified by the SVM.

Our SVM classification is based on the methods in Papageorgiou and Poggio (1999) since, of the 2D image analysis techniques, this method demonstrated the most encouraging results. Their single SVM was used to implicitly determine a model of pedestrian shape in a front/rear pose, whereas our system uses two SVMs, one to recognise

Figure 15. (a) Left image, (b) disparity map, (c) v-disparity image, (d) segmentation of the left image



pedestrians in a front/rear pose and another to recognise side pose. Therefore, segmented objects are classified with two SVMs.

The pedestrian shape was extracted by using a vertical and horizontal 3×3 Sobel edge detector. However, this pedestrian representation is too large to be efficiently handled by a SVM. Therefore, only the most distinguishing pixels from this edge image were chosen as the pedestrian representation. These pixels were chosen as those that consistently have high and low pixel values in the average edge images (Figure 16).

Figure 17 summarises the process of off-line SVM training and on-line SVM segment classification. Each SVM was trained on a database of approximately 1,500 positive and 20,000 negative grey scale images. Using an out-of-sample test image database (150 positive, 2000 negative images), our SVMs were capable of approximately 75% positive detection and 2% false detection.

Pedestrian Tracking

Both obstacle detection and classification generally provide robust results. However, the results can be incorrect, with obstacle localisation results being noisy and

Figure 16. Pedestrian shape representation (a) front, vertical edges, (b) front, horizontal edges, (c) side, vertical edges, (d) side, horizontal edges




Figure 17. Summary of object classification

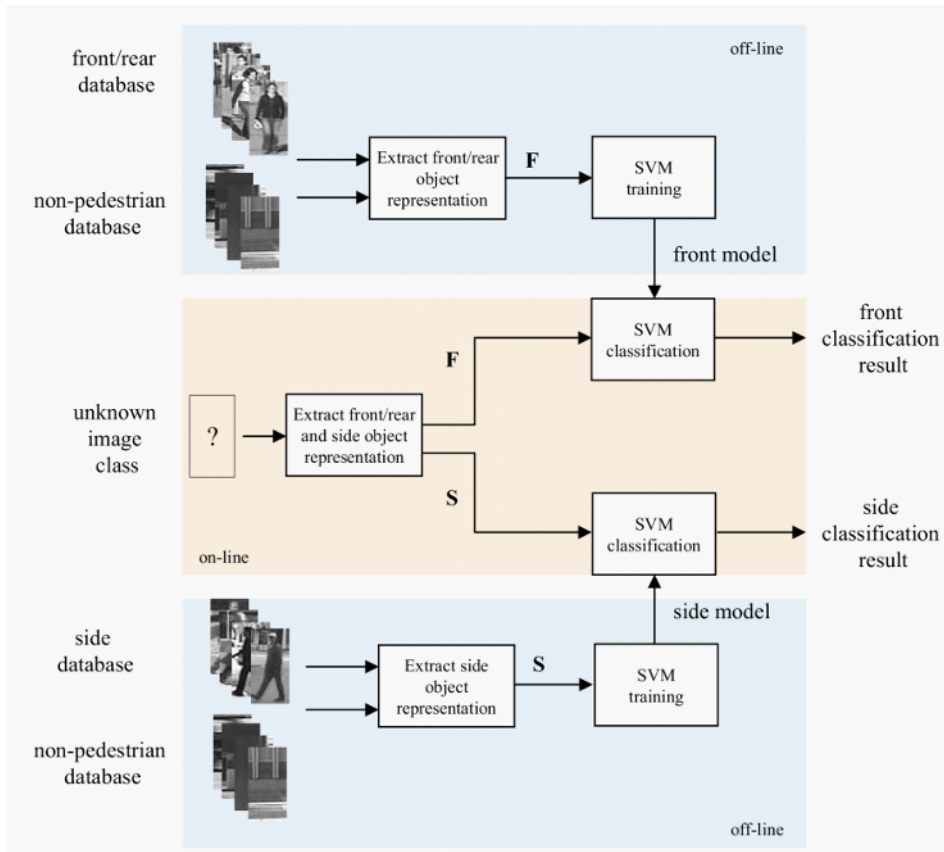
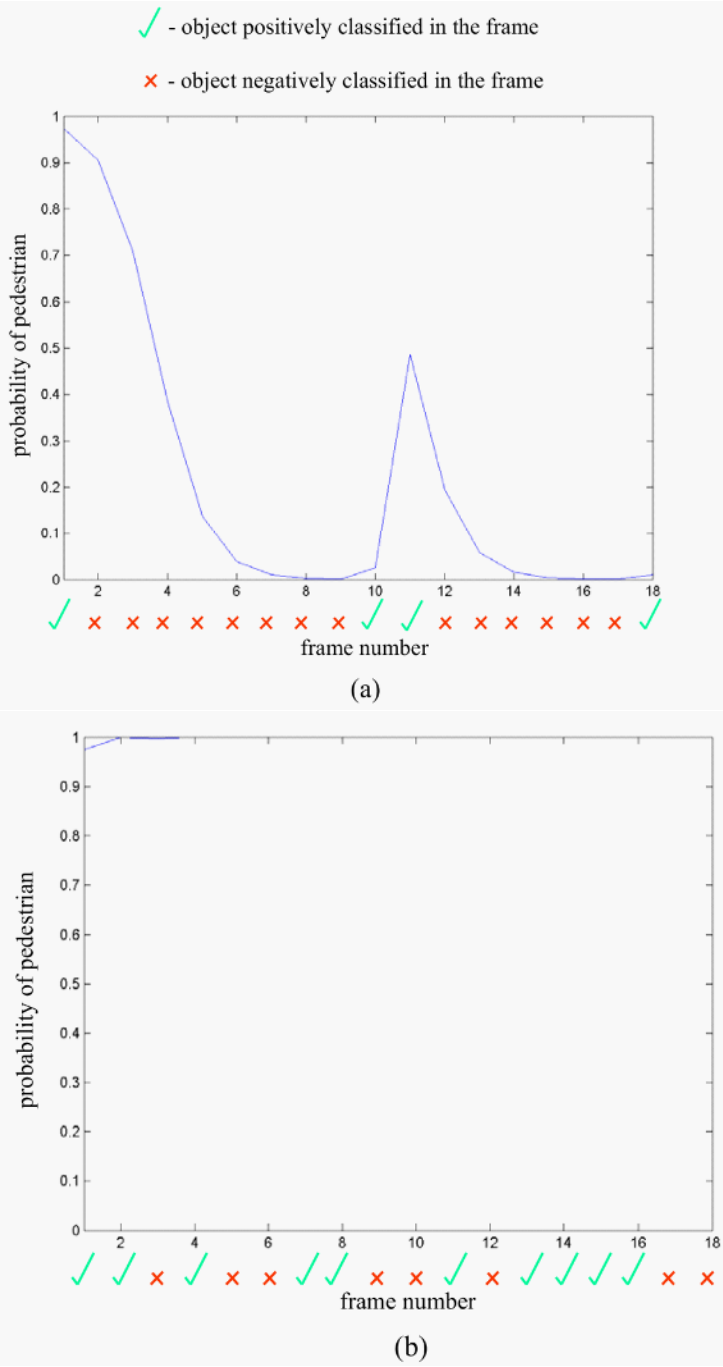


Figure 18. Example of how the probability that an object is a pedestrian varies with classification history (a) a typical non-pedestrian object, (b) a typical pedestrian



obstacle classification providing false detection. Pedestrian tracking aims to filter localisation results and minimise the effect of false positive and negative classification by temporally accumulating detection and classification estimates from each frame. The tracking algorithm combines a Kalman filter with a Bayesian approach to provide estimates of location, velocity, and pedestrian classification certainty over time.

Example results from analysing an object's classification history are depicted in Figure 18. These show how the probability of an object being a pedestrian varies with classification history. A pedestrian, typically receiving mostly positive classifications, has a certainty which remains high. On the other hand, a non-pedestrian, which typically receives few positive classifications, has a certainty which remains low.

Results

Our system was evaluated by in-vehicle testing in both simple and complex scenarios, with scene complexity rated according to 3D structure. Pedestrian tracking could be maintained with relative pedestrian speeds up to ~40km/hr (data processing power is the factor which limits tracking speed). Four scenarios were used to quantitatively determine detection rates. These sequences were analysed frame-by-frame. Table 1 presents a summary of results from the four sequences.

Figure 19 illustrates example results from the test sequences. On average, we achieved 83.5% positive detection and 0.4% false detection. The false detections are partly due to tracking continuing after the pedestrian left the field of view (caused by poorly selected tracking parameters). There were only very few obstacles tracked as a result of incorrect classification (in approximately 2,500 test frames analysed, only four non-pedestrians were momentarily tracked due to tracking being initiated from incorrect classification). This suggests a false detection rate of 0.3%.

To demonstrate the improved performance we achieved through using both 3D and temporal information, we modified our system to generate results based purely on a 2D image analysis and a 3D analysis with no temporal component. A comparison between the results obtained from such versions of our system is shown in Figure 20.

Table 1. Quantitative results from the four test sequences

Image Sequence Number	Scenario Description	Scenario Complexity	Number of critical pedestrians	Host Vehicle Speed (km/hr)	Positive Detection Rate (%)	Negative Detection Rate (%)
1	Through carpark	simple	1	15	98	0
2	Through carpark	moderate	3	20	75	0.2
3	Through busy pedestrian street	moderate	5	30	89	0.4
4	By congested pedestrian zone	difficult	5	20	80	0.2
Average					85.5	0.3

Figure 19. Example results from the four test scenarios

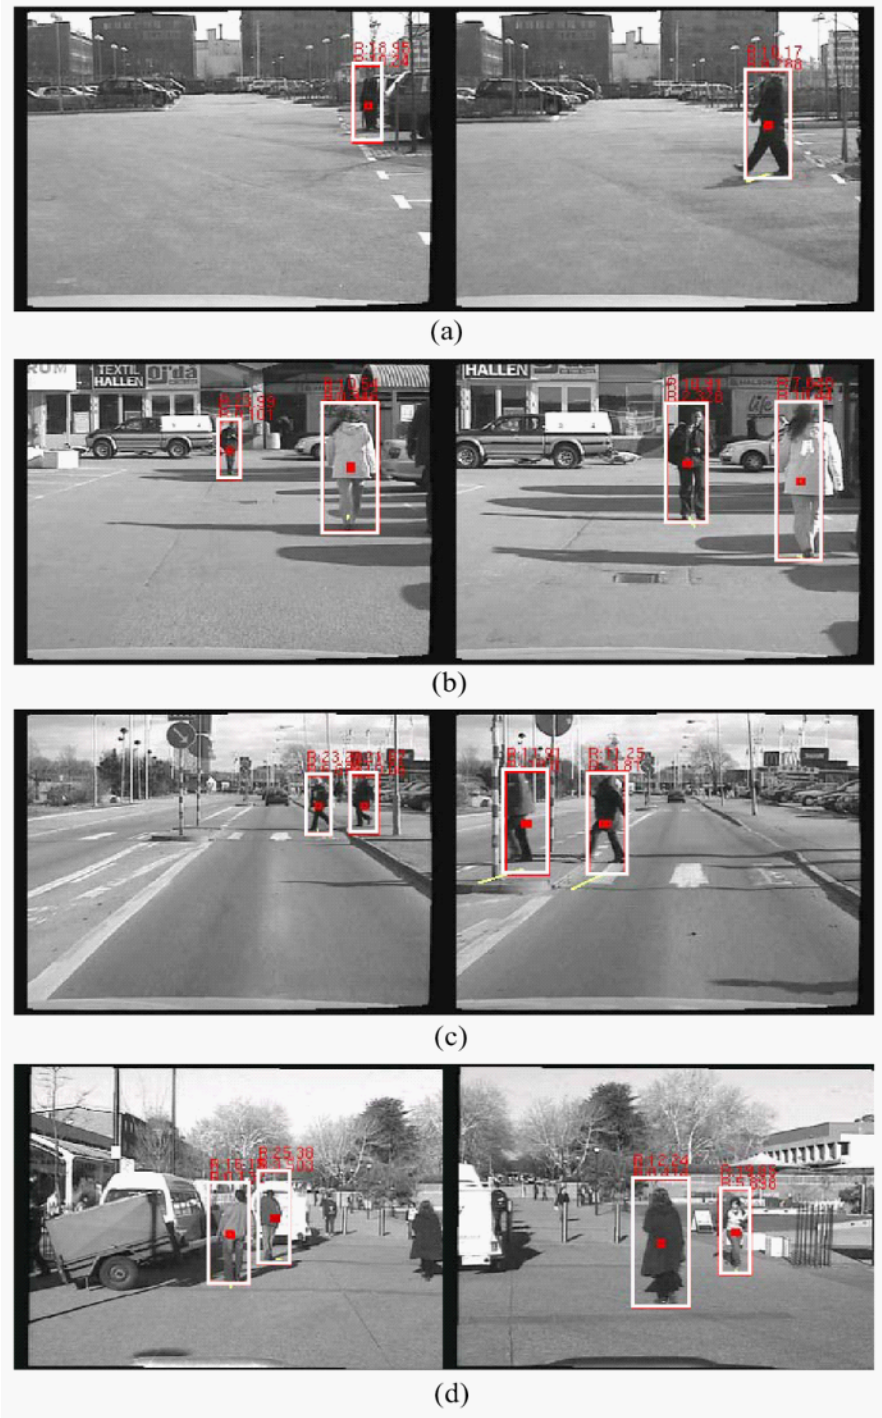




Figure 20. Comparison of results from different versions of our system to highlight the fact that a 3D temporal approach achieves superior results



Blind Spot Monitoring

Collisions can occur around automotive vehicles as a result of surrounding areas that are obscured from the driver's view point (Figure 21 (a)). In 2002, 58 children were backed over and killed in the U.S. by pickup trucks or sports utility vehicles, because the child was too small to be seen (Kids & Cars, 2003). Trucks have an added blind-spot area at the front corner of the vehicle, which is depicted in Figure 21(b). Blind spots are the cause of many accidents with other vehicles, as well as more vulnerable road users. The

Figure 21. (a) The area behind the car that is completely obscured from the driver's view point; (b) blind spot for a right hand drive truck



European Commission for Transport (2003) reports that in Belgium, an average of one cyclist is killed every month by turning trucks.

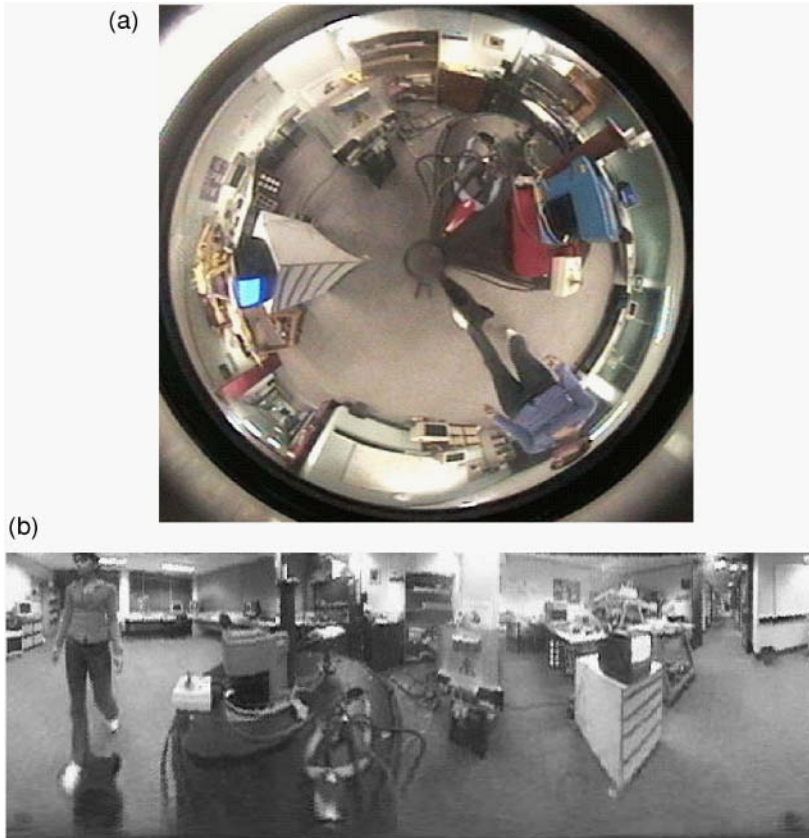
It is apparent from these statistics that there is a definite need to improve safety around cars and trucks, particularly concerning blind spots. There are several methods currently available to assist drivers in avoiding such collisions. The simplest method is the installation of devices to increase the driver's field of view, such as extra mirrors and wide-angle lenses. However, these still rely on an alert human observer. Another approach is to use an automated system using sensors such as sonar or radar. In these systems, range data is processed, and the driver is warned if a collision is imminent. The main drawback of the sonar sensors is their low angular resolution. A laser scanner could also be used, but only a thin two-dimensional plane can be monitored. Hence there is a need to develop new ways of sensing obstacles; one way is to use panoramic stereo vision.

Panoramic Imaging

Conventional cameras typically have one third or less of the perceived field of view of the human eye. The field of view on a camera can be increased in several ways. One method is to take many images from more than one static camera, or to use a single rotating camera (Szeliski, 1994). Disadvantages are the additional cost for more cameras, and the time to move and acquire one image for the latter. Furthermore, the need for moving parts decreases the robustness of the system. Alternately, a wide angle lens can provide a wide field of view without moving components, but such lenses are extremely bulky, expensive, and suffer from large angular distortion.

Convex mirrors are one approach to panoramic imaging that have been utilised extensively in the field of robotic navigation (Matsumoto, Ikeda, Inaba, & Inoue, 1999; Yagi, Kawato, & Tsuji, 1994). The sensor consists of a video camera which views a cone-like mirror. With the mirror on the optical axis, a full 360 can be viewed in the azimuth direction. The minimum and maximum angles of elevation captured are dependant upon the profile of the mirror surface. This method of panoramic imaging has several advantages. Since it is a passive sensor, its power requirements are small, and the lack of moving

Figure 22. (a) A raw image from a panoramic sensor which utilises a convex mirror; (b) the corresponding unwarped image



components means that the sensor could be made in a robust manner, which would require minimal maintenance.

An example of an image acquired by a panoramic sensor can be seen in Figure 22(a). These raw images are difficult for humans to understand, but they can be unwarped to create a more intuitive panorama, as seen in Figure 22(b). This allows the application of many conventional image processing techniques, which will be discussed further in the Obstacle Detection section.

The distortion introduced by the mirrors can be removed in several different ways. One method is to transform the panoramic image from a Cartesian to a polar coordinate system. Another approach is to project the warped image onto a virtual surface. In our system, we have chosen to project the image onto a virtual cylinder.

Stereo Panoramic Imaging

Binocular vision systems are used widely in computer vision for range estimation. However, these utilise cameras with small fields of view to minimise lens distortion. It is

also possible to estimate range using stereo panoramic systems. Conroy (2000) designed a stereo panoramic sensor that consisted of a single video camera and a double-lobed mirror. However, this system suffers from low image resolution and limited range-finding capabilities, due to the small distance (baseline) between the mirrors. In Gluckman, Nayar, and Thoresz (1998), Ollis, Herman, and Singh (1999), and Ng, Triverdi, and Ishiguro (1999), some theoretical and preliminary range estimation is investigated for systems of two or more panoramic sensors. Sogo and Ishiguro (2000) implemented a people-tracking system. However, the system was based on background subtraction, which is not easily extensible to this moving camera application.

System Overview

Our system consists of two panoramic sensors as shown in Figure 23(a). Approximately 200° of the sensor field of view is utilised.

Once the sensor has captured images of the blind-spot, these images are processed using an on-board PC to determine where obstacles are situated in the work space. The results can then be sent to a warning system to notify the driver, as shown in Figure 23(b).

Range Computations

The mirror profiles chosen do not conform to the single viewpoint constraint, and as a consequence, the range finding techniques used in conventional stereo vision would yield only approximate results. However, depth can be computed by using triangulation methods. Two corresponding points, P_1 and P_2 in the raw panoramic images, can be mapped back to a point on each mirror, using knowledge of the surface profile and the camera pinhole model. Using Snell’s law of reflection, it is then possible to determine the intersection of the two rays of light to estimate the position of a point P in three-dimensional space (Figure 24).

Sensor Resolution

The curved mirror surfaces create a more complicated depth resolution than for conventional stereo imaging systems, and this varies according to the mirror profile. The graph in Figure 25 was generated by computing the depth for every possible pair of image

Figure 23. (a) The stereo panoramic sensor, attached to the test vehicle in a horizontal configuration; (b) overview of the stereo panoramic vision system

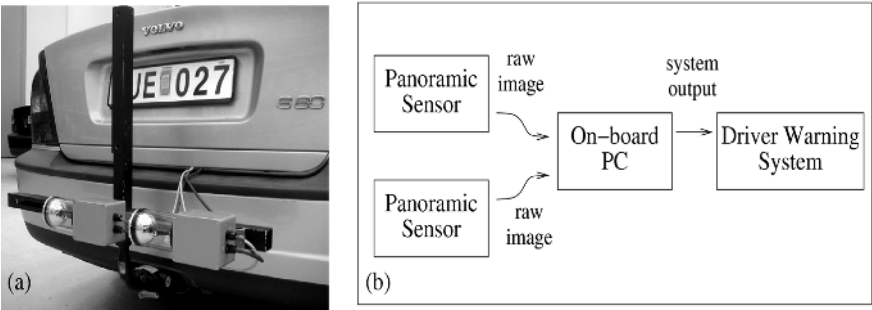
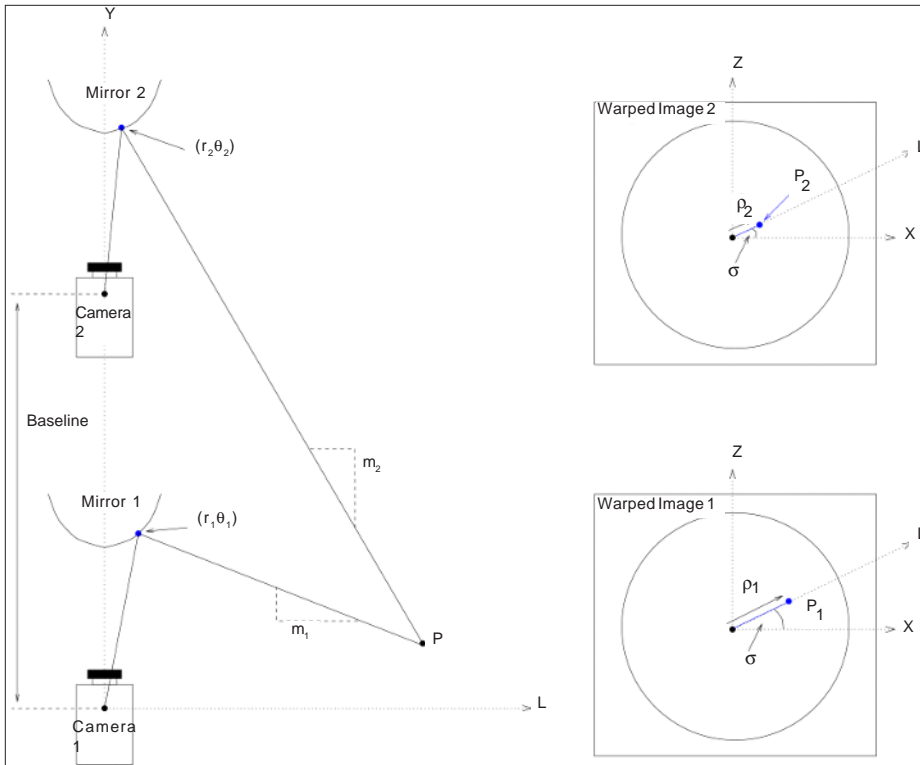


Figure 24. Range estimation of a point P in three dimensional space, given the projection P_1 and P_2 in the stereo panoramic images



pixel correspondences, up to a maximum disparity of 64. The depth was calculated using the approach described in the *Range Computations* Section. The maximum disparity search of 64 creates a circular *dead zone*, since objects in this area will have a disparity greater than this value. Such blind regions are, however, also present in the sensors mentioned in the *Blind Spot Monitoring* section.

Due to the limited CCD resolution, particularly toward the bottom of the panoramic image, discretisation errors become significant. The discretisation of three-dimensional space causes an error in the range and height estimation, because a point in space can only be mapped to the closest point in Figure 25. Figure 26(a) shows the maximum discretisation error along a ray of light with an elevation angle of 90° to the lower mirror axis. The discretisation error in height of an obstacle along this same ray of light is displayed in Figure 26(b).

Disparity Maps

With the camera axes aligned, the epipolar constraint corresponds to radial lines. When the images are unwarped, these become vertical parallel epipolar lines (Figure 27), and permit the application of many conventional image processing techniques. In this

Figure 25. Depth resolution of the stereo panoramic sensor, for a maximum disparity of 64 pixels


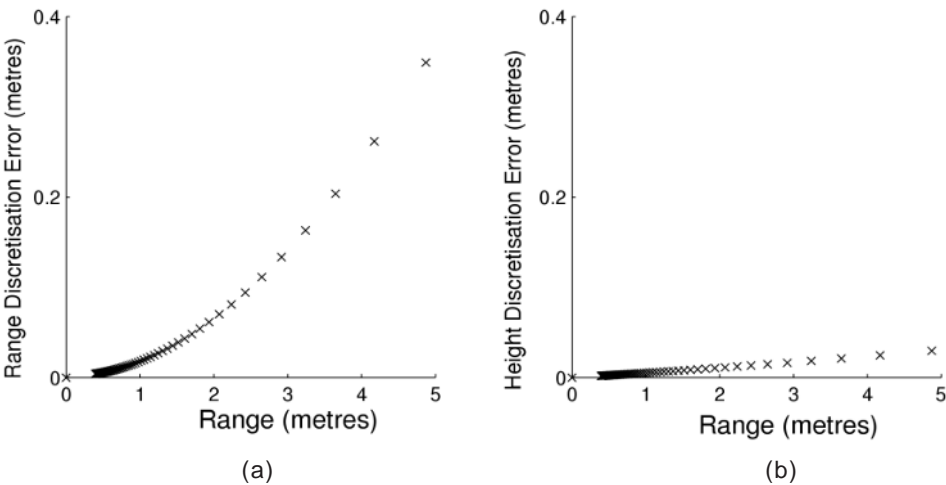


Figure 26. Discretisation errors of a typical panoramic sensor



case, disparity maps are generated by performing stereo matching along these lines using a standard window-based normalised cross correlation search.

Obstacle Detection

Obstacle detection was performed by first applying the v -disparity algorithm (Labayrade, Aubert, & Tarel, 2002) to the panoramic disparity maps and then segmenting the output. This algorithm was developed for conventional stereo vision systems, and

Figure 27. Epipolar lines are mapped from radial lines in the warped image (a) to parallel lines in the unwarped image (b)

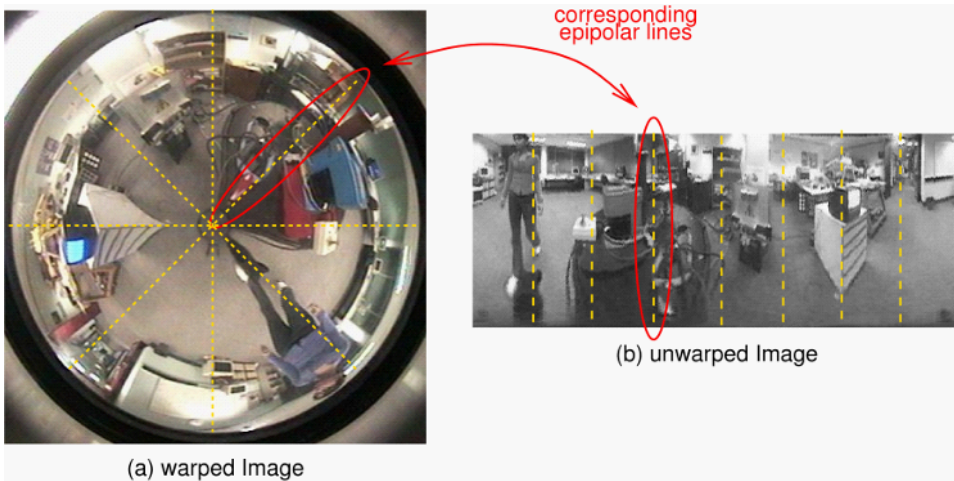
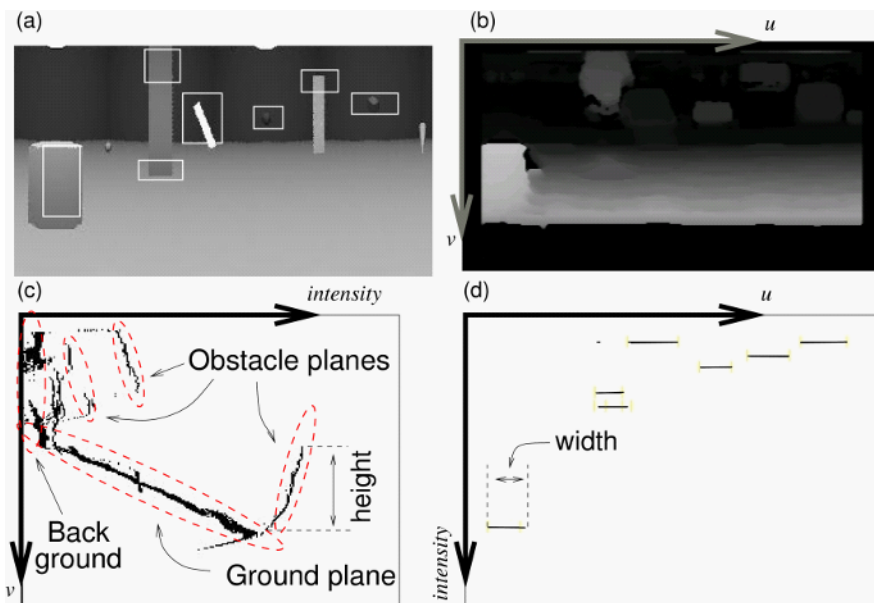


Figure 28. Obstacle detection results from the ray-traced images: (a) unwarped image, with obstacles detected, (b) disparity map, (c) v -disparity, (d) u -disparity



has not yet been utilised in panoramic vision. It is well suited to this application as it requires no *a priori* knowledge of the exact orientation of the ground plane, and is able to segment noisy disparity maps.

An example of a disparity map is displayed in Figure 28(b). The higher the intensity of a pixel, the higher the disparity and, therefore, the closer an object is to the stereo sensor. The v -disparity image is created by placing each pixel from the disparity map into bins according to their position along the vertical axis and intensity (Figure 28(c)). As a result, the ground plane appears as an angled line. Objects appear as near vertical lines above the ground plane, with the background displayed as the leftmost object. A Hough transform is used to detect the lines in the v -disparity image, and the height of the obstacles is determined by the length of the line. The width of an object is found by mapping these pixels to a u -disparity image, as shown in Figure 28(d). In this image, pixels are placed in bins according to their horizontal position and intensity value.


System Evaluation

The stereo panoramic system was evaluated to determine the range accuracy and effectiveness of the obstacle detection algorithm. This was done using ground truth data, and through field experiments.

A ground truth analysis was performed by creating artificial panoramic images using the Persistence of Vision (POV) Raytracer (<http://www.povray.org/>). Objects were placed in an environment around a virtual panoramic sensor at known locations, and the POV-Ray software was used to render images of this scene. Once the images had been unwarped, the corresponding features were selected manually, and range was calculated from these disparities.

As can be seen in Figure 29, the average range error generally remained below 20cm, with the standard deviation increasing to 76cm at a range of 6m. The increases in error

Figure 29. Results from the ground truth analysis (the average range error (left) generally remains within 0.2m; the standard deviation of the error is displayed on the right)



and standard deviation are partially due to human error; however, as expected, the majority of this can be attributed to the range discretisation error as shown in Figure 26. The average error in height estimation was much lower, always remaining below 15cm, while the standard deviation was no greater than 25cm. Again, this reflects the discretisation error in height estimation given in Figure 26. The estimation of the azimuth angular position was particularly impressive, with a maximum error of 3, and a standard deviation in the measurements of only 0.8° .

The obstacle detection algorithm was evaluated by automatically segmenting obstacles from the environment. Figure 28(a) displays results of the obstacle segmentation in the simulated environment. The algorithm was able to detect at least sections of all obstacles, except for one sphere, which did not have enough contrast with the background for the stereo matching algorithm to be successful.

Field Experiments

A three-by-five metre grid with one metre intervals was marked out on the ground. Images were captured and treated in the same way as described in the previous section.

In Figure 30, it can be seen that the error in range measurements was larger than that for the ground truth analysis. The error remained below 30cm until a range of 4m, where it increased. As well as human error and resolution issues, contributing factors are slight camera-mirror misalignment, and errors in the placement of the grid. However, the angular estimate was again very accurate, remaining within 2° of the actual position, increasing to within 5° at 4m; the standard deviation was always below 0.6° .

The disparity maps produced in these experiments were extremely noisy compared to those generated with ground truth data. Despite this, the v -disparity algorithm was able to produce results of a sufficient quality to successfully segment obstacles, as shown in Figure 31. However, due to the high noise ratio present in the real-world

Figure 30. The average range error (left), and range error standard deviation (right) calculated from the field experiments, using a sensor with a 31cm baseline (the estimated range generally falls within 0.4m of the true value)


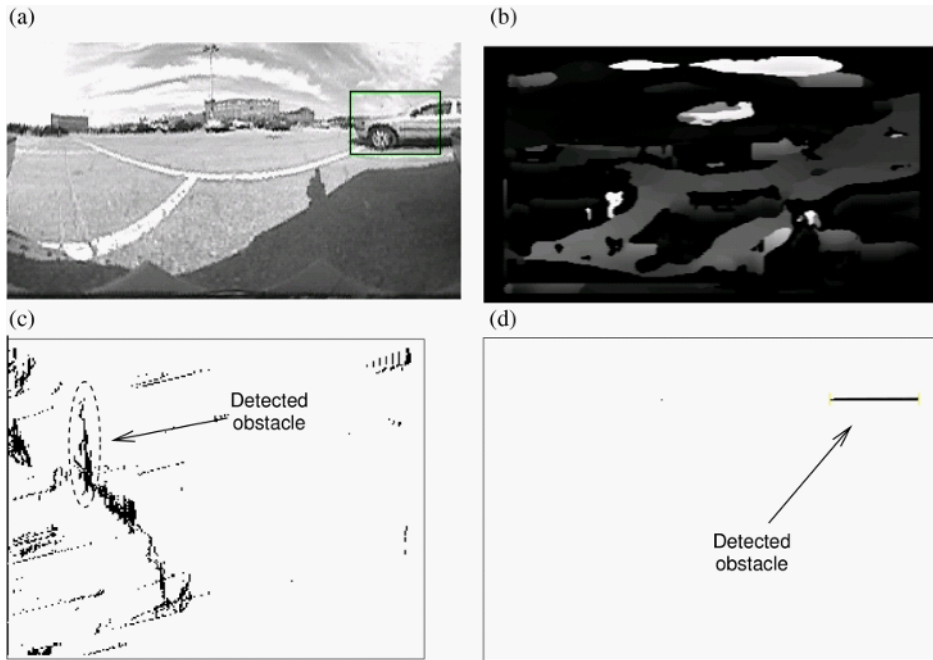


Figure 31. Obstacle detection results from the field experiments: (a) unwarped image, with obstacle detected. (b) disparity map. (c) v -disparity. (d) u -disparity



panoramic disparity images, false detection of obstacles became apparent in the image sequences. These generally only occurred in single frames and, as a result, false detections were easily filtered out by checking for temporal consistency. The system was modified to only report an object once it had been detected in at least two consecutive frames, and to continue to track the object until it had been lost in the same number of consecutive frames.

SUMMARY

We have suggested a new approach for obstacle detection, for the purpose of monitoring vehicle blind-spots. It was shown that stereo panoramic vision can be used to generate disparity maps from which objects can be segmented. This was done by applying the v -disparity algorithm, which has previously not been utilised in panoramic image processing. We found that this method was very powerful for segmenting obstacles, even in extremely noisy disparity maps. Our results indicate that range can be estimated reliably using a stereo panoramic sensor, with excellent angular accuracy in the azimuth direction. Furthermore, this sensor has the advantage of a much higher angular resolution and larger sensing volume than the driver assistance systems currently available.

CONCLUSION

In this chapter, we have given an overview of driver assistance systems in general together with a sample of the current research effort that is made within the Smart Cars project. The example systems presented show that today there are algorithms and techniques that are mature enough to be of practical use for driver assistance. Modern computers have a processing power that allows the use of advanced methods, increasing robustness, and reliability of the sensing algorithms. There is, however, a need to develop appropriate user interfaces. For example, if a pedestrian is detected to be on a collision course with the vehicle, what is then the best way to alert the driver without distracting him/her? Moreover, with the plethora of non-critical information that can be extracted from the road scene, how do we avoid overwhelming the driver? Human machine interfaces (HMI) is a research area that needs much attention in the future. An interesting area to also pursue is telematics used in vehicles. That is, exchanging information to and from the vehicle with road infrastructure, or other vehicles. An example of a telematics application may be, in the case of an accident, to automatically assess the status of the vehicle's passengers, and through wireless communication with the road infrastructure, provide information to rescue services.

ACKNOWLEDGMENTS

We would like to thank Leanne Matuszyk and Grant Grubb for their work on blind spot monitoring and pedestrian detection, which are parts of their Master Theses at the Australian National University in Canberra. Their work was also supported by Volvo Technology Corporation and Volvo Car Corporation. We would also like to gratefully acknowledge the support from National ICT Australia which now is responsible for the larger part of the Smart Cars project. National ICT Australia is funded by the Australian Government's Department of Communications, Information Technology, and the Arts, and the Australian Research Council through Backing Australia's Ability and the ICT Centre of Excellence programs.

REFERENCES

- Apostoloff, N., & Zelinsky, A. (2003). Vision in and out of vehicles: Integrated driver and road scene monitoring. *International Journal of Robotics Research*, 23(4-5), 513-538.
- Aufrere, R., Gowdy, J., Mertz, C., Thorpe, Wang, C. -C., & Yata, T. (2003). Perception for collision avoidance and autonomous driving. *Mechatronics*, 13(10), 1149-1161.
- Bertozzi, M., Broggi, A., Carletti, M., Fascioli, A., Graf, T., Grisleri, P., & Meinecke, M. (2003, September). IR pedestrian detection for advanced driver assistance systems. *Proceedings of the Pattern Recognition Symposium*, 2781, Magdeburg, Germany (pp. 582-590). Berlin: Springer-Verlag.
- Betke, M., & Makris, N. (1994). *Fast object recognition in noisy images using simulated annealing* (Tech. Rep. No. AIM-1510). Cambridge, MA: AI Laboratory, MIT.
- Broggi, A., Bertozzi, M., & Fascioli, A. (2001, May). Self-calibration of a stereo vision system for automotive applications. *Proceedings of the International Conference on Robotics and Automation*, Seoul, Korea (Vol. 4, pp. 3698-3702).

- Conroy, T. (2000). *Pulse train de-interleaving and panoramic vision systems*. PhD thesis, Research School of Information Science & Engineering, Australian National University.
- Dellaert, F., Pomerleau, D., & Thorpe, C. (1998, May). Model-based car tracking integrated with a road-follower. *Proceedings of the International Conference on Robotics and Automation*, Leuven, Belgium.
- Dickmanns, E. D. (1999, October). An expectation-based, multi-focal, saccadic (EMS) vision system for vehicle guidance. *Proceedings of the International Symposium on Robotics and Research*, Salt Lake City, Utah.
- Dickmanns, E. D. (2000, December). Vertebrate-type vision for autonomous vehicles. *Proceedings of the Symposium on Biologically Inspired Systems (BIS2000)*, Wollongong, Australia.
- Dickmanns, E. D., & Zapp, A. (1987, June). Automonous high speed road vehicle guidance by computer vision. *Proceedings of the International Federation of Automotive Control (IFAC) Conference*, Munich (pp. 232-237).
- European Commission. (n.d.). *Road Safety — Vehicles and Vehicle Equipment*. Retrieved July 2003, from http://europa.eu.int/comm/transport/road/index_en.htm
- Faugeras, O., Hotz, B., Mathieu, H., Viéville, T., Zhang, Z., Fua, P., Théron, E., Moll, L., Berry, G., Vuillemin, J., Bertin, P., & Proy, C. (1993). *Real-time correlation-based stereo: Algorithm, implementations, and applications* (Tech. Rep. No. RR-1013). France: INRIA.
- Franke, U., & Heinrich, S. (2002). Fast obstacle detection for urban traffic situations. *IEEE Transactions on Intelligent Transportation Systems*, 3(3), 173-181.
- Fredriksson, R., Haland, Y., & Yang, J. (2001, June). Evaluation of new pedestrian head injury protection system with a sensor in the bumper and lifting the bonnet's rear edge. *Proceedings of the International Technical Conference on The Enhanced Safety of Vehicles*, Amsterdam, The Netherlands (pp. 1-12).
- Fuerstenberg, K., & Lages, U. (2003, June). Pedestrian detection and classification by laserscanners. *Proceedings of the 9th EAEC International Congress*, Paris.
- Gavrila, D. (2000). Pedestrian detection from a moving vehicle. *Proceedings of the European Conference on Computer Vision*, Dublin, Ireland (Vol. 2, pp. 37-49).
- Gluckman, J., Nayar, S., & Thoresz, K. (1998, November). Real-time omnidirectional and panoramic stereo. *Proceedings of the DARPA Image Understanding Workshop*, Monterey, CA (pp. 299-303).
- Grover, R., Brooker, G., & Durrant-White, H. (2001, November). Low-level fusion of millimetre-wave radar and night vision imaging for enhanced characterisation of a cluttered environment. *Proceedings of the Australian Conference on Robotics and Automation (ACRA)*, Sydney, Australia (pp. 98-103).
- Johansson, B. (2002). *Road sign recognition from a moving vehicle*. Masters dissertation, Centre for Image Analysis, Swedish University of Agricultural Sciences.
- Kagami, S., Okada, K., Inaba, M., & Inoue, H. (2000, April 24-28). Design and implementation of onbody real-time depthmap generation system. *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco (pp. 1441-1446). Piscataway, NJ: IEEE Computer Press.
- Kids & Cars (2003, March). *Preliminary analysis of backing up accidents* (Tech. Rep.). Retrieved from www.kidsnadcars.org/legislation.html

- Labayrade, R., Aubert, D., & Tarel, J. P. (2002, June 17-21). Real-time obstacle detection in stereovision on non flat geometry through v-disparity representation. *Proceedings of the IEEE Intelligent Vehicle Symposium*, Versailles, France.
- Loy, G., & Zelinsky, A. (2003). Fast radial symmetry for detecting points of interest. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8), 959-973.
- Loy, G., Fletcher, L., Apostoloff, N., & Zelinsky, A. (2002, May). An adaptive fusion architecture for target tracking. *Proceedings of the 5th International Conference on Automatic Face and Gesture Recognition*, Washington, DC (pp. 261-266).
- Matsumoto, Y., Ikeda, K., Inaba, M., & Inoue, H. (1999, October). Visual navigation using omnidirectional view sequence. *Proceedings of the International Conference on Intelligent Robots and Systems*, Kyongju, Korea (pp. 317-322).
- Mertz, C., McNeil, S., & Thorpe, C. (2000, October 3-5). Side collision warning systems for transit buses. *Proceedings of the IEEE Intelligent Vehicle Symposium*, Dearborn, MI.
- Minor, L. G., & Sklansky, J. (1981). Detection and segmentation of blobs in infrared images. *IEEE Transactions on Systems, Man, and Cybernetics*, 11(3), 194-201.
- Miura, J., Kanda, T., & Shirai, Y. (2000, October 3-5). An active vision system for real-time traffic sign recognition. *Proceedings of the IEEE Intelligent Vehicles Symposium*, Dearborn, MI (pp. 52-57).
- Ng, K., Trivedi, M., & Ishiguro, H. (1999, January). 3D ranging and virtual view generation using omni-view cameras. In A. G. Tescher, B. Vasudev, V. M. Bove, Jr., & B. Derryberry, (Eds.), *Proceedings SPIE — Multimedia Systems and Applications*, 3528 (pp. 11-14).
- Ollis, M., Herman, H., & Singh, S. (1999, January). *Analysis and design of panoramic stereo vision using equi-angular pixel cameras* (Tech. Rep. No. CMU-RI-TR-99-04). Pittsburgh, PA: Carnegie Mellon University Robotics Institute.
- Paclik, P., Novovicova, J., Somol, P., & Pudil, P. (2000). Road sign classification using the laplace kernel classifier. *Pattern Recognition Letters*, 21, 1165-1173.
- Papageorgiou, C., & Poggio, T. (1999, October). Trainable pedestrian detection. *Proceedings of the International Conference on Image Processing*, Kobe, Japan (pp. 35-39).
- Piccioli, G., De Micheli, E., Parodi, P., & Campani, M. (1996). Robust method for road sign detection and recognition. *Image and Vision Computing*, 14(3), 209-223.
- Priese, L., Klieber, J., Lakmann, R., Rehmann, V., & Schian, R. (1994, August). New results on traffic sign recognition. *Proceedings of the Intelligent Vehicles Symposium*, Paris (pp. 249-254). Piscataway, NJ: IEEE Press.
- Roberts, J., & Corke, P. (2000, August). Obstacle detection for a mining vehicle using a 2D laser. *Proceedings of the Australian Conference on Robotics and Automation (ACRA)*, Melbourne, Australia (pp. 185-190).
- Simoncelli, E. P. (1999). *Bayesian multi-scale differential optical flow*. San Diego, CA: Academic Press.
- Sogo, T., & Ishiguro, H. (2000, June). Real-time target localization and tracking by n-ocular stereo. *Proceedings of the IEEE Workshop on Omnidirectional Vision*, Hilton Head, SC (p. 153).
- Szeliski, R. (1994, December). Image mosaicing for tele-reality applications. *Proceedings of the IEEE Workshop on Applications of Computer Vision*, Sarasota, FL (pp. 44-53).

- Yagi, Y., Kawato, S., & Tsuji, S. (1994). Real-time omnidirectional image sensor (COPIS) for vision-guided navigation. *IEEE Transactions on Robotics and Automation*, 10, 11-21.
- Zhao, L., & Thorpe, C. (2000). Stereo and neural network-based pedestrian detection. *IEEE Transactions on Intelligent Transportation Systems*, 1(3), 148 -154.

Chapter VI

The Application of Swarm Intelligence to Collective Robots

Amanda J. C. Sharkey, University of Sheffield, UK

Noel Sharkey, University of Sheffield, UK

ABSTRACT

This chapter considers the application of swarm intelligence principles to collective robotics. Our aim is to identify the reasons for the growing interest in the intersection of these two areas, and to evaluate the progress that has been made to date. In the course of this chapter, we will discuss the implications of taking a swarm intelligent approach, and review recent research and applications. The area of “swarm robotics” offers considerable promise for practical application, although it is still in its infancy, and many of the tasks that have been achieved are better described as “proof-of-concept” examples, rather than full-blown applications. In the first part of the chapter, we will examine what taking a swarm intelligence approach to robotics implies, and outline its expected benefits. We shall then proceed to review recent swarm robotic applications, before concluding with a case study application of predator-prey robotics that illustrates some of the potential of the approach.

TAKING A SWARM INTELLIGENCE APPROACH TO COLLECTIVE ROBOTICS

What is Swarm Intelligence?

To be able to deliberate over the reasons for taking a swarm intelligence approach to collective robotics, we need to first provide an account of what swarm intelligence is. Swarm intelligence is a comparatively recently articulated notion. The concept was first introduced by Beni and Wang (1989) in their investigations of simulated self-organising agents in the context of cellular robotic systems. A more extensive definition was provided by Bonabeau, Dorigo, and Theraulaz (1999), who suggested the term should be applied to:

any attempt to design algorithms or distributed problem-solving devices inspired by the collective behaviour of social insect colonies ... and other animal societies.

Agassounon, Martinoli, and Easton (2004) suggest that swarm intelligence takes its inspiration:

from the biological examples provided by social insects ... such as ants, termites, bees, and wasps, and by swarming, flocking, herding, and shoaling phenomena in vertebrates.

By contrast, Martinoli (1999) asserts that such bio-inspiration is not necessary, and that the defining characteristic of swarm intelligence should be an emphasis on local control and communication (as opposed to global), claiming that:

swarm intelligence arises from local interactions and is based on local information and communication mechanisms. (Martinoli, 1999)

Nonetheless, an emphasis on local as opposed to global interaction is itself clearly biologically-inspired. Bonabeau and Theraulaz (2000) suggest the term “swarm intelligence” is applicable to the “collective behaviour that emerges from a group of social insects”. It would seem then, that the term can be applied to both the emergent collective behaviour of biological swarms or colonies, *and* to algorithms inspired by living systems.

In either case, the notion of swarm intelligence is grounded in an awareness of the sophisticated collective behaviour that can emerge from the combination of many simple individuals, each operating autonomously. Despite their often extensive size — colonies of the African driver ant *Anomma wilverthi* may contain as many as 22 million workers patrolling an area as much as 50,000 square meters in extent (Raignier & van Boven, 1955) — insect societies are able to maintain themselves as a collective, and to accomplish the coordinated action needed to construct nests, to feed and raise their young, and to react to invasion or other interference despite, or perhaps because of, the limited behavioural and representational capabilities of their individual members, and the absence of centralised control mechanisms.

Swarm intelligence algorithms have been shown to be useful for both static problems (for example, application of the ant colony system to the *travelling salesperson problem* — Dorigo & Gambardella, 1997), and dynamic problems, notably load balancing in telecommunication networks (Schoonderwoerd, Holland, Bruten, & Rothkrantz, 1997). In the ant colony system's solution for the travelling salesperson problem, a set of agents, or "ants" search for good solutions and communicate through pheromone-mediated indirect communication. The results show the system to be competitive with other heuristic algorithms such as genetic algorithms, evolutionary programming, and simulated annealing. Interestingly, although the system is biologically inspired, its operation departs from that of real ants through the introduction of faster pheromone decay. In real ants, if a shorter path is presented after a longer path, it is not adopted because the longer path will have been marked by pheromone; whereas in an artificial system, this problem is avoided by introducing pheromone decay (Bonabeau & Theraulaz, 2000). The work of Schoonderwoerd et al. (1997) also relies in the application of an ant colony algorithm, and on the removal of obsolete solutions by applying a mathematical version of pheromone evaporation.

A swarm intelligence approach has then been applied to a number of tasks and practical applications. The notion, abstracted away from biology, depends on decentralised local control of a large number of simple agents. The role of the environment is stressed, although that environment is often virtual rather than real. Although a swarm intelligence system will contain no explicit model of the environment, individual agents can both receive information about the environment and act on that environment to change it. The advantages of swarm intelligence include the idea that the resulting collective systems will be scalable because the same control architecture is used for both a few and for thousands of units. This results in increased flexibility because the individual units can be dynamically added or removed without the need for explicit reorganisation. It also increases robustness because of the reliance on unit redundancy and minimalist unit design. At the same time the reliance on autonomy and self-sufficiency can increase the flexibility of the system and its ability to adapt quickly to rapidly-changing situations.

Swarm Robotics

Our concern here is with the application of swarm intelligence principles to collective robotics. The result can be termed "swarm robotics", although there is some debate about the definition of this term. Nonetheless, the emerging consensus seems to be that swarm robotics involves groups of simple robots: that are autonomous; that are not controlled centrally or remotely; that are capable only of local communication; and whose operation is in some sense biologically inspired. Such robot collections inherit many of the advantages of swarm intelligence, and can be seen to offer a number of benefits if deployed in environments that are unknown, hostile, and liable to unpredictable change. We shall consider these advantages in more detail below, but can summarise them briefly here. The absence of centralised control can reduce problems associated with communication bottlenecks and delays. Similarly, the insistence on local rather than global communication, or in most cases, the absence of any inter-robot communication, means that as in swarm intelligence, the same control system can be used for any number of robots without the need for recalibration and adjustments to communication proto-

cols. The use of many cheap and expendable robots with limited communicative and sensorial abilities can also reduce costs. At the same time, a swarm intelligence approach can profit from reliability through redundancy, since the ability to complete a task need not be jeopardised by the failure of individual robots, or of a centralised controller. In addition, as argued by Brooks in founding papers in this area, (Brooks, 1999), autonomous simple robots can react to the environment in real time, without the need to complete intensive sense-plan-act cycles of analysis. And the stress in autonomous robotics on a close relationship to, and exploitation of, the physical environment, particularly when combined with adaptive techniques for developing robot control systems, can often have the result of finding simpler and more effective solutions than those that are usually obtained by more traditional methods that depend on human designers.

In the following sub-sections, we shall consider some of the key characteristics of swarm robotics in turn: biological inspiration; individual simplicity; control; communication; and finally, group composition. In our considerations, we will address centrality of that characteristic, and the benefits and limitations likely to result from it.

Biological Inspiration

A close relationship to biology is one of the defining features, and one of the strengths, of taking a swarm intelligence approach to robot collectives. The notion of biological inspiration is almost explicit in the term “swarm intelligence”, and is central to both swarm robotics and behaviour-based robotics. Animals, as Sharkey (2003) comments, “*exhibit ...remarkable capacities for flexible adaptation to novel circumstances*”, and a strength of modern robotics is its ability to capture some of the capabilities of biological entities. Biologically-inspired robotics has its roots in the seminal work of Grey Walter (1953), and re-emerged in the 1980’s with Brooks’ development of behaviour-based robotics (Brooks, 1999), and Braitenberg’s demonstrations of the emergence of complex behaviours from a combination of very simple neural networks encoding different taxes (Braitenberg, 1984).

The relationship between behaviour-based robotics and biological systems is a reciprocal one: an understanding of the mechanisms underlying animal and insect behaviour can be exploited to build flexible robotic control systems; and such systems can be used to test the behavioural consequences of biological models. Sharkey (2003) identifies generalised and specific classes of bio-robotics research: generalised research being a strand dedicated to applying broad notions from the life sciences to robot control, whilst specific research is geared towards testing the more specific implications of biological models. In the same article, Sharkey identifies a third, theoretical strand in which the implications of such bio-robotics research for theories and views of life, mind, and cognition, are explored, a strand of limited relevance to more practical concerns and one that is not explored further here.

Generalised research, in which our understanding of the biological mechanisms that underlie animal and insect behaviour is exploited, is clearly related to our present concerns. Such research can involve, for example, exploiting an understanding of the simple mechanisms that underlie the seemingly complex behaviours of social insects such as bees and ants, to develop biologically-inspired control systems for robots. For instance, some of the organisational mechanisms of ants or bees could be adapted for collections of robots, to arrive at novel emergent solutions (see section below on *Group*

Composition). Likewise, models of bird flocking, fish schooling, and toad detour behaviour have formed the basis of reactive robotic systems (Arkin, 2003). And the adoption of evolutionary methods, and the research area of evolutionary robotics can also be seen as examples of generalised bio-robotics, as are learning techniques such as reinforcement learning (see the *Explicitly Cooperative Tasks* section on methods for developing control systems).

Specific biologically-inspired research is also of some, albeit more indirect, relevance to our present concerns. It is only possible to apply mechanisms based on those of biological systems if those systems themselves are understood. Hence, detailed and specific accounts of biological systems (for example, of ants cooperating to move large prey items — see references in Kube & Bonabeau, 2000) are relevant, as are papers in which explanations of the underlying mechanisms are tested. And such research can directly lead to the development of solutions for practical tasks to which collections of robots can be applied. For example, all of these elements can be found in a study by Kube and Bonabeau (2000) (further described below in the section on *Applications*). In their paper, they provide a model of cooperative transport in ants, from which they derive testable predictions about the kind of stagnation recovery mechanisms to be expected depending on ecological conditions, and prey size. Kube and Bonabeau suggest that, “because the model is able to reproduce many of the collective features of cooperative transport in ants with a minimum of plausible assumptions, it suggests that these assumptions may be sufficient to explain the behaviour observed in ants”. At the same time, the repositioning and realigning behaviour observed in ants when copied in a robotic system provides a solution to the stagnation that results when several robots apply equal force to push an object from several opposing directions at once.

In brief, biological inspiration is a central tenet of swarm robotics and behaviour-based robotics. The whole notion of swarm intelligence depends on such inspiration. What is less clear is how strictly such inspiration should be interpreted: there is the minimalist position, according to which the abilities of robots in a collective should be restricted to a minimal level in order to gain the advantages of a swarm intelligent approach; and there is also the position that some of the specialist capabilities afforded by modern electronics should be exploited where possible. In the main, the swarm robotic studies described in the next section, and in the case study, are better described as minimalist.

Individual Simplicity

The simplicity of the individual robots in a collective is another of the features emphasised in a swarm intelligence-based approach, as it is in behaviour-based robotics. The use of simple robots can result in a less costly system, and one that is more robust in the sense that simple robots are less likely to fail (since there is less to go wrong). Another advantage of simplicity is the ability to respond rapidly and flexibly to changes in the environment. The kind of simplicity we have in mind refers in particular to the control system used. Individual robots in a collective can themselves be subject to reactive, or deliberative control. A *reactive* architecture is one that “tightly couples perception to action without the use of intervening abstract representations” (Arkin, 1998). A *deliberative* architecture, on the other hand, relies on abstract representations of the world. Between the two lie those systems that extend purely reactive systems with

some memory capabilities; in other words, rather than just reacting to a stimulus, the robots are affected by an internal register which has some form of memory.

The earlier formulations of behaviour-based robotics stressed the importance of reactivity (Brooks, 1986, 1991), and its advantages in terms of the ability to respond quickly to stimuli in the environment. More recent formulations of behaviour-based robotics have incorporated some degree of memory and representation needed to accomplish more complex tasks, while maintaining an emphasis on avoiding the use of centralised representations and control (Mataric, 1997b). Arkin (1990) advocated the use of control systems consisting of a hybrid of reactive and deliberative control. However, the situations in which deliberative control is likely to be the preferred option are those in which uncertainty is limited, and the world can be accurately modelled, not the kinds of situation for which swarm robotics are best suited. An emphasis on the application of swarm intelligence principles to collective robotics implies the use of control systems that are as simple and reactive as possible.

The simplest control system for an individual robot, then, is one in which control is as close as possible to sensors and actuators as, for example, when an artificial neural network is used and the inputs are stimulated by raw sensor values, while the outputs control motor speed and direction. At a higher level, the architecture can be organised into basic behaviours, each representing a perception-action loop. The starting point is usually behavioural modules responsible for robot movement (for example, a *wander* module, and an *obstacle avoidance* module). Higher level modules responsible for finding objects or moving towards a goal can then be added, depending on the task in question. Decisions about the design of behavioural modules are usually the responsibility of a human designer. Once a set of behavioural modules is chosen, some method of combining them is required. Behavioural modules can be combined by means of either a *selection* method (switching control to the most appropriate module), or by *fusing* them. We shall consider examples of each of these in turn.

Probably the best known selection method is that represented by Brooks' subsumption architecture. In a subsumption architecture (Brooks, 1986), a fixed priority scheme is defined for basic behaviours such that enabling one of them results in the suppression or inhibition of others, so that only one behaviour is active at any one time. An alternative switching mechanism was proposed by Maes (1989) based on spreading activation between modules. A more recent development of Maes' system was proposed by Jung and Zelinsky (1999): a selection method termed architecture for behaviour-based agents, (ABBA). The selection mechanism is based on a winner-take-all scheme. Activation is spread among competence modules on the basis of the output of feature detectors, and the pre-conditioning competence module. When the activation level of a competence module reaches threshold, it becomes active. The scheme has been tested on two heterogeneous Yamabico robots (<http://www.roboken.esys.tsukuba.ac.jp/english/Yamabico>) performing a collaborative cleaning task.

Switching and selection methods rely on the assumption that only one behavioural module should be active at any one time. The alternative approach is to adopt some form of fusion of modules, where the outputs of several active modules are fused, or combined in some way, to result in a single behaviour that reflects the influence of several modules. For example, under the motor schema-based approach (Arkin, 1989), primitive behaviours, or motor schemes, can be active simultaneously, and combined cooperatively. Behaviour

is obtained by multiplying the vector response of each motor schema by a gain, and then summing and normalizing the result. The DAMN architecture (Payton, Rosenblatt, & Keirsey, 1990) used by researchers at Carnegie-Mellon University for controlling unmanned ground vehicles, similarly relies on fusion, using a scheme by which each behaviour votes for and against each of a set of possible vehicle actions, and an arbiter performs command fusion to select the most appropriate action.

To summarise: The application of swarm intelligence to collective robotics implies the need for simple robots that can respond rapidly and flexibly to the environment. The main way to achieve this, at present, is to rely on a system of reactive control at the individual level, or a set of reactive behavioural modules combined through some form of action selection.

Collective Control

We now turn to a consideration of the way in which a collection of robots can be coordinated in order to accomplish a task. The emphasis in swarm intelligence is on decentralised control, or autonomy. In robotics, a collection of autonomous robots is an example of decentralised control, since there is no centralised controller responsible for their coordination. Insect societies are similarly only locally, or indirectly controlled; there is no central body or agent that issues commands to organise the nest or colony. One of the advantages of decentralised control in swarm robotics is an increase in fault tolerance (again, there is no risk that a centralised controller will fail and result in a deterioration or breakdown of the system). Individual autonomous robots can also respond more quickly and flexibly to a changing environment, since they can respond directly to information from their own sensors, and do not need to wait for centralised instructions.

It is possible for cooperative behaviour to emerge as the result of the combined effect of individual behaviours. Few would disagree that the cooperation found in insect societies is the result of emergent properties, rather than planning. Similarly, instances of apparently cooperative behaviour can be found in collections of autonomous robots; examples are described in the following section on applications. A classic example of decentralised control of a group of robots is that of the Frisbee-sorting robots of Holland and Melhuish (1999). Here the cooperation that occurs is emergent, since the individual robots are simple and autonomous, and incapable of direct communication with each other: they each follow a fixed set of reactive rules.

Questions about how best to achieve such emergent behaviours in robots are currently an active focus of research. One method is to handcraft the rules or control mechanisms. The adaptive, cooperative behaviour of the Frisbee-sorting robots is the combined result of individual robots following a fixed set of reactive rules, and shows that it is quite possible to generate adaptive behaviour without the use of learning algorithms. There are many cases of adaptive behaviours that are genetically determined: for example, “hard-wired” reflexes and instincts. Control mechanisms for robots could then, for example, consist of relatively simple “rule-like” mechanisms, encoded in terms of handcrafted weights for a neural network, such that when an obstacle is detected via a robot’s sensors, the result is that it turns away from it. Similarly, a subsumption-based system of behavioural modules could result in adaptive behaviour, despite having been handcrafted rather than learnt or evolved. The *Case Study* section provides an example.

A popular and effective alternative to handcrafting such control mechanisms is to use genetic algorithms to evolve them. Nolfi and Floreano (2000), in their book on evolutionary robotics, describe some of the numerous examples of the elegant solutions that can be obtained when evolutionary techniques allow the environment to determine their design. Neural networks are particularly useful control mechanisms for autonomous robots because of their flexible response to noisy environments, but it is not obvious how to provide detailed training feedback about what a robot should do at each time step. An evolutionary algorithm, and its associated fitness function, provides a mechanism for an overall evaluation of the performance of the network over the entire evaluation period. It can also be used to evolve *any* parameter of the neural network, in other words, not just the weights, but also the learning rule, neuron transfer function, and network topology. Some interesting work has also been initiated in which evolutionary techniques have been applied to the modular design of a neural network controller (Nolfi, 1997), although this approach has not yet been fully developed or exploited.

Evolutionary robotics has been shown to be useful in the development of effective control mechanisms for individual robots. The feedback it provides also offers a mechanism for developing effective control strategies for a multi-robot team, although research in this area is very much in its infancy. To date, questions about the evolution of multi-robot, or multi-agent systems have primarily been investigated in simulation (for instance, Baldassarre, Nolfi, & Parisi, 2002; Martinoli, 1999), but we anticipate that this is an area that will receive an increasing amount of attention in the near future.

Currently, questions about how to use evolutionary techniques to evolve robotic control structures form a particularly active focus for research. Of course, it is never the case that a control system is simply evolved in response to the environment: The researcher has always made some contribution. Choices and decisions will have had to be made about certain aspects such as the environment, the task, the initial architecture, and evolutionary operators such as the fitness function, for example. Nonetheless there is still much to be said for keeping such intervention as much to a minimum as possible, based on the underlying idea of getting as close as possible to emulating natural evolution. An interesting study in terms of an attempt to reduce the level of experimenter intervention is a technique termed “embodied evolution” (Watson, Ficici, & Pollack, 2002). Watson, Ficici, and Pollack’s reported aim is to create an evolutionary technique that can be run automatically on a group of robots without the need for global communication, or a reliance on simulation. Robots are left in the task environment (the task in question being one of phototaxis, or attempting to reach a light source from different starting positions). The evolutionary mechanism used was crossover, with a simple neural network control architecture serving as the evolutionary substrate. Using a mechanism termed probabilistic gene transfer algorithm (PGTA), each robot maintains a virtual energy level that reflects its performance, and probabilistically broadcasts genetic information locally, at a rate proportional to its energy level. Robots that are close enough to each other will pick up this information, and allow the broadcast genes to overwrite some of their own. Robots accept broadcast genes with a probability inversely related to their own energy levels. The result is that those robots with a higher energy level (because they have performed the task more effectively) are more able to broadcast information, and less likely to allow their own genes to be overwritten. The method, when tested, compared favourably to a hand-designed solution to the same task.

There is also interest in the combination of evolutionary and learning techniques, where a neural network control structure is further refined by a lifetime learning process (Nolfi, 2003). Such an approach can be used to develop individuals with a predisposition to learn: evolving effective starting conditions, or initial weight matrices; or evolving the tendency to behave in such a way that the individual is exposed to appropriate learning experiences. The consequence of combining evolution and learning has been shown in some studies (for example, Nolfi & Parisi, 1997) to lead to some promising results. A further area that offers promise is to explore the application of evolutionary techniques to the design decisions about the initial set of behavioural modules and to the methods used to combine them.

An interesting alternative to evolving neural network weights for control systems is to develop a training set by remotely controlling a robot around an environment, and collecting examples of inputs and corresponding appropriate motor responses (Sharkey, 1998). Other than this, the main alternative to handcrafting neural networks, or evolving their weights, is to use reinforcement learning algorithms. Such algorithms have the advantage of not requiring a training set; what is needed instead is a scalar evaluation of behaviour to guide learning. The evaluation could be provided by a trainer, or by the agent itself as a result of its interaction with the environment. What is needed is a policy, or mapping from states to actions, that maximises positive reinforcement. Various algorithms have been explored from Q-learning (Watkins, 1989) to the learning classifier system advocated by Dorigo and Colombetti (1998) that incorporates a reinforcement algorithm closely related to Q-learning to adjust the strength of its rules, and a genetic algorithm to search the space of possible rules.

Reinforcement learning has also been applied to the issue of collective behaviour. Mataric (1997) presents a method through which four mobile robots learn, through reinforcement, social rules about yielding and sharing information in a foraging task. The problem is one of finding a way of reinforcing individual robots for behaviour that benefits the group as a whole, since greedy individualist strategies will result in poor performance in group situations with resource competition. The solution she investigates is one that relies on social reinforcement for appropriate use of four behaviours: yielding, proceeding, communicating, and listening. Her results indicate improved foraging behaviour in a group of four mobile robots subject to social reinforcement, as compared to a group using only greedy individual strategies. In the social reinforcement condition, robots were rewarded (a) for making progress to sub-goals of finding food, or returning home (b) for repeating another robot's behaviour and (c) for observing reinforcement delivered to another robot; in other words, vicarious reinforcement, whether positive or negative, is shared amongst all robots in the locality.

In conclusion: The application of swarm intelligent principles to the control of robot collection is best exemplified by studies in which no use is made of global control. Decentralised control is needed to achieve the full advantages of scalability and redundancy, and there is a considerable body of research that has investigated different methods of achieving desired collective behaviours without resorting to global control. There are also some other examples, such as that of ALLIANCE architecture (Parker, 1998), that share several features with swarm robotics (for example, the use of robots that individually are subject to decentralised control), but in which some form of global control, or global communication (see next sub-section) is employed.

Communication

Communication is of course closely linked to the control of robot collectives. Again our concern here is with swarm intelligence, and hence with minimal and local communication. This is an area that is beginning to receive more attention. For some in the area such as Luc Steels (1996), the concern is to develop explanations of the evolution of complex communicative abilities. For others, the concern is to make use of limited communicative abilities to extend the capabilities of the group. Communication of some form is likely to be required to accomplish some form of task allocation, and the coordination needed to jointly undertake and complete a task. Clearly biological systems such as those formed by social insects depend on some forms of communication. These include alarm, recruitment, and recognition (Wilson, 1971), but also indirect communication by means of the environment. Cao, Fukunaga, and Kahng (1997) identify three major forms of inter-robot communication: interaction via (a) the environment (b) sensing and (c) communications. We shall make use of their distinctions, since they are useful in terms of robots, but note that the distinction between sensing and communication is not a particularly useful one in terms of social insects, since the communication they are capable of depends on their ability to sense chemical pheromones.

Interaction via the Environment

The simplest form of interaction relies on communication through the environment, a form of communication that is clearly biologically inspired. The term *stigmergy* is used to refer to such communication, and was proposed first by Grassé (1959) in the context of his studies of termite nest building. In termites' stigmergic labour, it is the product of work previously accomplished, not direct communication between nest mates, which induces insects to perform further labour. For example, Grassé distinguished three successive stages in the construction of a single foundation arch by workers of *Macrotermes (Bellicositermes) bellicosus*. In the first stage, when workers encounter building material of pellets of soil and excrement in a container, they all explore the container individually. In a subsequent stage of "uncoordinated work" the pellets are carried about and put down in a seemingly haphazard fashion. Eventually, seemingly by chance, two or three pellets get stuck on top and workers begin to add more pellets on top, until a column begins to grow. If a second column is located nearby, when a certain height is reached, they begin to bend the column towards the neighbouring column, with the result that eventually the two columns meet, the arch is finished, and the workers move away.

Although Grassé's claim that the stigmergic explanation is perfect has been criticised (for instance, Stuart, 1967, pointed out the inability of a simple stigmergic machine to shut down when the job is finished), the concept of indirect communication via the environment remains a useful one in swarm robotics. Holland and Melhuish's (1999) study of sorting examines the operation of stigmergy and self-organisation in a homogenous group of physical robots. Their robots move and drop Frisbees, or pucks, in an arena, in a manner that depends on their encounters with pucks that have been dropped by other robots. If a puck is encountered by a robot that is not already carrying one, the robot picks it up in its gripper. If a robot bumps into a puck when it is already carrying one in its gripper, this causes the robot to release its gripper and to drop the puck it was carrying. The effect of following such simple behavioural rules is that the pucks

eventually end up being clustered together: a result that can be described as stigmergic, since the robots are responding to the “work” of other robots that have deposited the Frisbees. Further rules such as responding slightly differently to pucks of two different colours, and reversing backwards before dropping pucks of one colour can be shown to result in sorting the pucks. The task is claimed to be analogous to the brood sorting behaviours of ants, and similar rules seem to describe the ants’ behaviour. The study provides a clear demonstration of the exploitation of real-world physics, and of the emergence of higher-level behaviour (sorting) from the combined effect of a number of robots following simple reactive rules.

Interaction via Sensing

Interaction via sensing refers to local interactions that occur between robots as a result of their ability to sense one another, but without explicit communication. Such sensing would permit a robot to distinguish between other robots, and objects in the environment, sometimes termed “kin recognition” (Mataric, 1993). A recent example of this would be the “infrared sniffing” implemented by Noel Sharkey on the robots showcased at Magna (see *Case Study* section), where the predators can detect the prey and vice versa, by detecting the infrared signals they emit, while still relying on sending out infrared signals themselves in order to detect other obstacles and objects in the environment. Social insects are also clearly able to recognise nest mates, an ability made clear by their concomitant ability to detect alien insect intruders. They can also be shown to be able to distinguish between castes and life stages of nest mates (Wilson, 1971), and it seems likely that chemical pheromones are implicated in these abilities.

Interaction via Communication

There are some systems, such as the ALLIANCE architecture (Parker, 1998) mentioned in the preceding sub-section on *Collective Control*, that rely on global communication, allowing each robot to be aware of the activities of all other team mates. Parker argues that this is legitimised by the poor sensory capabilities available in autonomous robotics. However, such communication does not fit well with the biologically-inspired notion of swarm intelligence. In addition, it is possible to identify a number of disadvantages to a reliance on global communication, as outlined by Martinoli (1999). First, global communication is difficult to scale up, and bottlenecks soon arise when group size increases. Second, such communication requires more sophisticated and expensive hardware, which can reduce the robustness, and increase the expense of the robots. And third, clearly social insects are not aware of all the activities of other colony members, but are able to coordinate their activities. They rely instead on stigmergic signals and local communication. It makes sense, therefore, where simple and expendable robots are appropriate, to rely instead on simpler forms of communication.

Nonetheless, some robotic applications such as those of Parker (1998) do rely on global communication. There is also some research, exemplified by that of Steels (1996) in which linguistic communication between robots is explored, with the aim of exploring the evolutionary emergence of language. This research is of interest from the point of view of cognitive science, and of exploring evolutionary explanations of the origin of language, but is of limited relevance to more practical domains.

Even though insect societies are not capable of global communication whereby the same information is broadcast to all the members of the society, they are capable of local

communication. They predominantly make use of chemical signals, as opposed to visual or auditory signals. They constantly touch each other, but do not seem to make use of touch to convey much information. Chemical signals, or pheromones, are implicated (at least) in alarm and recruitment, and in recognition according to Wilson (1971). Even in the classic example used by Grassé to explain nest building in termites, it is likely (Wilson, 1971) that more than stigmergy is involved. Stuart (1967) looked at the conditions that affected the halting of repair work to termite nest walls. He found that chemical communication was involved, as odour trails were used to recruit workers to the scene when a breach in the nest wall occurred. Termites continued building, or repairing, until the disturbing stimuli of air currents and lowered humidity were removed.

To date there has been little attempt in swarm robotics to incorporate mechanisms of local communication; the emphasis has rather been on minimalism, and keeping the robots as simple as possible. Clearly, however, the equivalent of pheromone trails are made use of in swarm intelligence applications (see the *What is Swarm Intelligence?* section), and are beginning to be explored in robotic studies also. It is quite possible to justify local communication on the basis of the source of biological inspiration for the approach of swarm robotics; since ants and bees use it, it may well turn out to be useful in swarm robotics also, and it is quite likely that future developments will find ways to make greater use of such communication in future.

Group Composition

Initially, swarm robotic research, such as Holland and Melhuish (1999), focused on the development of swarms of identical robots. More recently, there has been an increased interest in the deployment of different types of robot. In their taxonomy of multi-agent systems, Dudek, Jenkin and Miliotis (2002) include group composition (homogeneous or heterogeneous) as one of the axes that can be used to discriminate between collectives. It should, however, be pointed out that even a group of seemingly identical robots will become heterogeneous, as differences in sensor tuning, calibration, robot drift, and wear and tear amplify initially negligible differences (in the *Case Study* section, it is found that certain Predator robots, are much more effective at catching Prey than others, even though their software and hardware are intended to be identical). Nonetheless, while acknowledging that homogeneity-heterogeneity is best viewed as a continuum rather than a discrete classification, it remains the case that some studies are explicitly concerned with a group of robots that can at least initially be considered to be identical (homogenous groups), while others are concerned with groups that are clearly not identical since they differ in their mechanics, their sensing, their role within the group, or their controllers, or underlying basic behaviours.

The development and use of heterogeneous groups of robots can be justified both in terms of biological inspiration and practical applications. We can briefly consider the biological justification first. The social organisation of ants, social bees and wasps, and termites all depends on polymorphism, defined as the co-existence of two or more functionally-different castes within the same sex. Three basic female castes, for instance, can be found in ants: the worker, the soldier, and the queen. Termites also have a soldier caste specialised for colony defence, and a worker caste. The existence of castes provides a mechanism for the division of labour that depends both on morphological differences and on the age of the insects, since there is usually a temporal division of

labour in a sequence that leads from nest-work, to brood care, to foraging. Even a brief consideration of the organisation of social insects makes it clear that current investigations in swarm robotics, while biologically inspired, do not even come close to reflecting their complexity and effectiveness.

Nonetheless, there are clearly advantages to the use of heterogeneous groups of robots for real-world applications. Often an application demands capabilities that cannot be easily built into a single robot: a robot cannot be both big and small at the same time; similarly, a single robot may not be able to carry all the sensors needed for a particular task. For example, Grabowski, Navarro-Serment, Panedis, and Khosla, (2002) describe a heterogeneous team they developed for mapping and exploration: a team that consists of four types of robot (large All Terrain Vehicles, medium sized Tank robots, Pioneer robots, and centimetre scale Millibots). The All Terrain Vehicles can transport the smaller robots to distant places of interest; the Pioneers are designed to facilitate exchange of information between team members; the Tank robots are autonomous and can undertake individual missions or coordinate the Millibots; and the Millibots are so small that they can manoeuvre into small spaces. Their team is hierarchically organised, and applied in the task domain of exploration and mapping.

Another example of a heterogeneous group is the marsupial robots employed by Robin Murphy and her colleagues. Murphy (2002) argues for the benefits of transporter or marsupial robots, enabling the transportation of small task-specific robots to the target area, without loss of battery power. The domain she is particularly interested in is that of search and rescue, and in her marsupial robot teams, *mother* robots can transport *daughter* robots over rubble to the target site, and can also offer backup and protection, for instance, recharging facilities, collection and processing of sensor data, and shelter from environmental conditions (such as planetary nights).

It can be seen then that there are a number of practical advantages to forming a collection of robots of varied abilities. Also, from the biological point of view, as Birk and Belpaeme (1998) point out, ecosystems with only one species are not biologically plausible. The many issues to be investigated in this area include: the measurement and representation of the degree of heterogeneity; task allocation between members of a heterogeneous team; physical cooperation and coordination between members; and communication.

Issues about the measurement and representation of heterogeneity have been explored by a number of researchers. The most obvious way in which robots in a heterogeneous team might differ is in terms of their physical form. For example, a group of robots might differ in its method of locomotion, which of course affects their mobility. For instance, in the marsupial robots investigated by Murphy, the mothers and daughters are physically different in that the mothers have the ability to transport and protect the daughters. Grabowski et al. (2002) also consider the effect of different forms of propulsion in their heterogeneous groups. Kephra robots, for example, are wheeled, with small wheels housed in the centre of the robot (<http://www.k-team.com/robots/>). This is good for flat surfaces, but not for inclines. The millibots they use, on the other hand, can be equipped with thick rubber treads, allowing them to climb inclines. A heterogeneous team could also differ in their sensorial capabilities, and in behavioural capabilities. Other researchers have explored heterogeneity in groups of robots that exists only in software; the heterogeneous group of robots studied by Ijspeert, Martinoli, Billard, and Gambardella, (2001) differ only in the length of time they grip the sticks, in a stick pulling experiment.

An important issue in heterogeneous and in homogenous robot teams is the way in which task allocation is carried out. One way in which this can be accomplished is through the adoption of specialised roles. For example, in the research described by Goldberg and Mataric (2002), the number of collisions between robots is reduced by using a dominance hierarchy; dominant robots are given priority. An alternative method is to use individual activation thresholds for task allocation. Although it used to be assumed that task allocation within insect societies was a rigid process (Gordon, 1996), more recent research has focused on behavioural flexibility and stressed the importance of external and decentralised factors such as pheromones or individual encounters (for example, Bourke & Franks, 1995). In an activation-threshold model, individuals react to stimuli intrinsically bound to the task in question. For example, neglected brood, or the corpses of dead ants, diffuse an odour of increasing strength. When this stimulus reaches threshold value, an individual reacts by performing the relevant activity (such as grooming the brood, or carrying a corpse out of the nest). If individuals do not have the same threshold values, recruitment is gradual, and team size is thereby regulated. Krieger and Billeter (2000), using teams of up to 12 real robots, implemented a simple and decentralised task allocation mechanism based on individual activation thresholds. Their results show that this mechanism resulted in efficient and dynamical task allocation.

Another related issue is that of finding a method of measuring the degree of heterogeneity present in a group. Parker (1994), in her PhD thesis, introduced the concept of *task coverage*, which measures the ability of a given team member to achieve a specific task, a measure which decreases as groups become more heterogeneous. Balch (2002) introduces a method for measuring robot group diversity, based on *social entropy*, and argues for the importance of a quantitative metric. He concentrates on evaluating diversity in teams of mechanically-similar agents that use reinforcement learning to develop behavioural policies.

In summary, the issue of group composition, and the move from collections of seemingly identical robots to the development of heterogeneous groups, is one that is increasingly coming to the fore in collective robotics. Interestingly, it is not something that seems to be considered in swarm intelligence research, despite its biological justification. However, in swarm robotics there seems to be a gradual movement towards increasing task differentiation, and heterogeneity as more complex applications are considered and attempted.

APPLICATIONS

The task domains for which collections of simple autonomous robots seem most appropriate are those that occur in areas that are inaccessible or hazardous to humans, and that are likely to benefit from having a number of small, light, expendable, and cheap robots. These include surveillance, monitoring, de-mining (detecting and removing mines), toxic waste disposal, exploration, and search. However, in the current state of research, such tasks are usually not actually carried out, but rather analogous tasks, or tasks that involve components of these, are investigated. Foraging, for instance, involves many subcomponents of the task associated with toxic waste clean up, while

being closely related to the foraging behaviour of biological agents such as ants. Kube and Bonabeau (2000) in a brief review of tasks to which a swarm intelligence approach has been taken, write:

As the reader will perhaps be disappointed by the simplicity of the tasks performed by state-of-the-art swarm-based robotic systems let us remind her or him.. (that) .. it seems urgent to work at the fundamental level of what algorithms should be put into these robots: understanding the nature of coordination in groups of simple agents is a first step towards implementing useful multirobot systems.

In the following sub-sections, we will look at a number of task domains in which swarm robotic, or collective robotic, solutions have been sought, although as explained, these often explore only components of an application. Where real-world applications have been developed (for instance, in Urban Search and Rescue), these often involve a compromise in which swarm intelligent methods are combined with global control and communication, or even remote control and teleoperation.

Traffic Control and Moving in Formation

As soon as collections of robots are used, issues of traffic control become relevant. Cao, Fukunaga, and Kahng (1997) in their review of collective robotics, report the use of traffic rules by several researchers to avoid collision and deadlock. Goldberg and Mataric (2002) describe a pack controller that reduces the number of collisions between robots returning home, by establishing a dominance hierarchy. Robots returning home are given priority depending on their position in the dominance hierarchy, an approach shown to significantly reduce the number of collisions compared to a homogenous robotic controller. This and related solutions are of interest, although the need for traffic control is reduced in swarm robotic studies because of the emphasis in behaviour-based robotics on collision avoidance as one of the most basic behaviours.

Moving in formation is usually investigated as a leader-follower task in which one robot follows another. Desai, Ostrowski, and Kumar (1998) report a study in which robots can communicate locally, and where by means of such communication they converge and move together. Dudek, Jenkin, Milios, and Wilkes, (1996) describe leader-follower experiments in which a leader robot signals its intention to the follower robot by making specific motions prior to the intended movement.

Maintenance of Energy: Recharging and Artificial Ecosystems

The ability of simple robots to locate and return to a base for recharging when necessary is one that would be particularly useful when they are deployed in inaccessible areas. The maintenance of energy has been researched in a few studies, although often this is simulated rather than actual. Grey Walter's research (1953) provided an early example of a recharging task. His electronic "tortoise" was attracted to and docked with a recharging station when its power was low. When it had sufficient energy, it was

repelled by the recharger. The behaviour of the tortoise was driven by simple reactive behaviours: head toward weak light, back away from strong light, and turn-and-push to avoid obstacles. The case study below provides a further example of an artificial ecosystem, as does the study reported by Birk and Belpaeme (1998), in which robots competed with energy-draining lamps for power.

Explicitly-Cooperative Tasks

Some tasks or scenarios can benefit *indirectly* from cooperation between robots (such as predator robots showing emergent cooperation to catch prey robots). However, there are also some tasks that are *explicitly* designed to require cooperation. Box pushing, where the box in question is too heavy to be pushed by a single individual, is a case in point. Similarly, a task of pulling sticks out of the ground can be set up so as to require cooperation.

Box Pushing

Parker (2002) considers the application of a learning version of ALLIANCE to a box-pushing task, using the task to address issues raised by removing and replacing team members. However, as mentioned earlier, the ALLIANCE architecture depends on global communicative abilities that are not representative of swarm intelligence. Kube and Zhang (1996) describe a box-pushing system in which a number of robots move boxes to an indicated goal. Kube and Bonabeau (2000) further explore the same task and its relationship to the cooperative transport performed by ants. Interestingly, some of the behaviours observed in ants are adapted for use in the robots. When cooperatively transporting prey, ants show realigning and repositioning behaviours. These same behaviours were adapted for use with robots, to overcome the stagnation that can result from several robots pushing the object in different directions. The robotic system relies on the robots' tendency to move towards the brightly-lit goal. Stagnation is avoided through the use of some of the repositioning and realigning strategies noticed in ants.

Robotic Soccer

Robotic soccer is an explicitly-cooperative game by definition, but relies on a fixed size of team, and usually on global observation and communication to players. Nonetheless, it provides an interesting framework within which to investigate cooperative behaviours, although much of this has been done in simulation (for example, Pagello, D'Angelo, Montesello, Garelli, & Ferrari, 1999).

Cooperative Multi-Robot Observation of Multiple Moving Targets

This is a domain investigated by Parker and her colleagues that they describe as an inherently cooperative task, and a rich test bed for research on multi-robot cooperation, learning, and adaptation. Parker (2002) considers methods for applying reinforcement learning techniques to the problem, showing that it is possible to come close to a human-generated solution.

Stick Pulling

This task is one that is set up so as to require the collaboration of two robots to be successful. The task is to locate sticks in a circular arena and to pull them out of the ground, where the length of the stick means that a single robot cannot pull it out of the ground alone, but must collaborate with a second robot. Ijspeert et al. (2001) report a study in which the effect on task performance of the number of robots and of gripping time is investigated. The study is of particular interest since it provides one of the few examples of a collaborative robotic experiment where the robots are controlled using swarm intelligence principles, and they are physical, not simulated robots, and more than three robots (up to six were used). Their robots are based on swarm intelligence principles; they depend on reactive control and minimal sensing capabilities. Their results, using real robots, a web simulator, and a probabilistic model, show that collaboration can result in the absence of signalling or planning, or communication other than a stigmergic communication in which the state of the environment is changed. They varied the number of robots, the number of sticks, and the gripping time parameters for individual robots. Collaboration rates increased as more robots were added to a group. They found that a heterogeneous group of robots using a simple signalling scheme (heterogeneity being introduced at the software level, with robots differing from each other in their gripping time parameter) increased collaboration rate under some circumstances.

Foraging

Essentially, in foraging, as, for example, interpreted by Goldberg and Mataric (2002), robots search designated regions of space for certain objects, and once found, bring the objects to a pre-specified goal region. As such, the task is closely related to toxic waste cleanup. It also has relevance to search and rescue (see below), and is one of the canonical test beds for cooperative robotics (Cao, Fukunaga, & Kahng 1997). A variety of approaches have been taken to the task, ranging from simple stigmergy (Beckers, Holland, & Deneubourg, 1994), to the formation of chains along which objects are passed to the goal (Drgoul & Feber, 1993).

Exploration, Mapping, and Search

Although most of the exploration algorithms that have been proposed have dealt with single robot exploration, a number of multi-robot exploration algorithms have begun to appear in the literature. Given the time and cost associated with exploration, it seems a particularly suitable task for a robot collective. Yamauchi (1999) investigates a method for the decentralised coordination of multi-robot exploration. Autonomously-controlled robots explore using frontier-based exploration. Every robot maintains its own map, an occupancy grid, but communicates it also to other robots.

Each robot stores the local grids received from other robots and integrates them with its own local grid to form its own global map reflecting what is known about the territory. The approach as described is interesting in the way in which robots both share information, but also maintain their own maps and make independent decisions about where to explore. The advantages of an occupancy, or evidence grid, lie in the way in which they can be used to fuse information from different types of sensors and from

different robots. A related approach in which different guesses about the location of a robot are combined in an evidence grid, is pursued in research by Gerecke, Sharkey, and Sharkey (2000). The communication of maps between robots means that the method is less biologically plausible than swarm-based methods depending on local communication; however, the shared communication of honey-bees as they dance to indicate the location of nectar sources, or potential nest-sites, bears some similarities.

Other examples of multi-robot exploration exist. Burgard, Moors, Fox, Simmons, and Thrun (2000) provide another example of two autonomous robots cooperating to construct a probability-based occupancy grid representation of space. Dudek, Jenkin, & Milios (1996) show that a collective of robots can explore a topological (graph-like) environment more effectively than a single robot. Robots independently explore parts of the graph, keeping track of where they have visited, and meeting on a prearranged schedule to merge their maps and subdivide the remaining unexplored portions of the graph. Rekleitis, Dudek, and Milos (1997) show how two or more robots can coordinate a motion strategy to construct an accurate metric map even without odometry.

Grabowski et al. (2002) describe a mapping and exploration application, in which small millibots explore an environment and communicate it to a team leader, who collects all sensor information and uses an occupancy grid representation to build up a composite map of the area. However, the system relies on keeping track of the location of all the robots, using an ultrasound-based localization system in which millibots act as both beacons and localization receivers.

Coordinated search has also been investigated. Spires and Goldsmith (1998) raise the question of how best to coordinate a group of robots to efficiently search an area for targets (such as landmines). The solution they suggest, and for which they present some preliminary results, is a decentralised approach in which robots traverse a space-filling curve.

In all these examples, the aim is to exploit the increased information obtaining capability of a group, as opposed to a single robot, and to find efficient ways of combining this information. The exploitation of the increased information-gathering capabilities of several redundant robots can be seen to have parallels in the current enthusiasm in the machine learning community for ensemble approaches (Sharkey, 1999), in which the outputs of several redundant classifiers can be combined to form a more accurate classification.

Search and Rescue

Urban search and rescue is a real application domain in which collections of robots have been deployed. Robin Murphy's team at the Centre for Robot-Assisted Search and Rescue, (CRASAR) University of South Florida, is apparently ready to respond to national and international events within four hours. They participated in the search and rescue at the September 11th World Trade Towers collapse, deploying robots at the scene within six hours. They used "man packable" robots which could be carried in backpacks by one or two people. Their emphasis is on the development of software for agile robots that can access areas that are inaccessible to humans. As with most practical applications, the approach taken is not a purist one; in the field, a combination of human remote control and a hybrid deliberative/reactive control mechanism is used.

Predator-Prey Scenarios

Research on heterogeneous groups has not always looked at cooperation. Various researchers have investigated competition in predator/prey scenarios. Nolfi and Floreano (2000), for instance, have considered the way in which the evolution of two competing populations with coupled fitness may reciprocally drive each other to increasing levels of complexity. Most such investigations have considered only a single example of both predator and prey. By contrast, Noel Sharkey's predator and prey robots (see *Case Study* section below) at the Magna Science Adventure Centre in Rotherham consist of several predators, and several prey: a set up that facilitates the investigation of the cooperation that can emerge as a result of applying evolutionary mechanisms to groups of robots. Predator-prey scenarios can thus be used to explore both competitive and cooperative relationships between robots. Another related example is that of Birk and Belpaeme (1998) who describe a multi-agent system also based on heterogeneous robots. They developed an artificial ecosystem containing three different species, and describe the potential for cooperation and competition that is present in the scenario.

CASE STUDY: PREDATOR AND PREY ROBOTS AT MAGNA

In this penultimate section of our chapter, we describe the predator-and-prey robots developed by one of the authors (Noel Sharkey) as a museum exhibit for public engagement with science at the Magna Science Adventure Centre in Rotherham, UK (<http://www.visitmagna.co.uk/>). In the earlier *Maintenance of Energy* section, we introduced the concept of artificial ecologies, as pioneered by Grey Walter. Sharkey extended the biological inspiration of recharging to the development of an artificial food chain where the main currency was electricity. Some of the robots recharge under light

Figure 1. Predator robot



Figure 2. Prey robot



Figure 3. Prey robot “feeding” in light



and other robots steal the charge — as an analogy to the food chain in natural systems. The robot colony is divided into two distinct “species” consisting of five *Predators* (Figure 1) and ten *Prey* (Figure 2).

To maintain energy levels, the *Prey* have solar panels on their tops to recharge their batteries by driving them under light trees (Figure 3) which they must first seek out using infrared and solar sensors.

Figure 4. Predator robot catching prey



Figure 5. Predator robot spiking prey



The *Predators*, on the other hand, must hunt the *Prey* to maintain their energy levels. They capture the *Prey* (Figure 4) by lifting them off the ground and then they extract energy by driving a metal fang onto a brass plate in the middle of the *Prey* (Figures 5 and 6).

The task of the prey was to escape capture from the predators and maintain energy levels by charging under light trees, while the task of the predators was to catch the prey and steal battery power.

Figure 6. Close-up of prey top

One of the main aims was to illustrate the effective behaviour that can arise in a collection of reactive robots equipped with very simple sensing capabilities. Each of the robots was equipped with eight *infrared* (IR) sensors. These were equally distributed around the circular prey — the dark “windows” just beneath the top in Figure 2. On the oval-shaped predators, there were two infrared (IR) sensors on each side and one on the back. Because of the hunting task, the predator had three sensors focused on the front. The only other sensors used were *bump sensors* and *solar sensors*. Each predator had two bump sensors on the front to detect contact with the prey. These were inset so that they would only trigger on contact with a prey robot and not a wall. The prey had solar sensors to enable centring under lights.

The IR sensors have three different functions for the colony. First, the intensity of the reflected IR light is used to detect distances to objects for navigating around the arena without bumping into obstacles and other robots. Second, the IRs are used on the prey robots for detecting bright lights for recharging.

Finally, the IRs are the only means that the robots have for detecting the presence of other robots. Each robot species emits a unique IR signature in the form of pulses. The width of these pulses determines if a robot is “friend” or “foe” with the predator pulse about 2.5 times wider than the prey pulse. This novel use was termed, “infrared sniffing”. The prey are sensitive to the infrared signature of the predators and use it to detect their approach. The predators use the infrared signature of the prey to detect and hunt them.

The robots are controlled by a number of independent behaviour modules that are initially prioritised and then triggered by environmental cues — similar to Brook’s (1986) subsumption architecture. The behaviour of the prey is controlled by six modules:

1. **Object module** — an 8x2 artificial neural network that takes input from the IR distance-sensing mode and steers the robot away from objects.

2. **Predator sensing module** — an 8x2 artificial neural network (ANN) that takes input from the “infrared sniffing” mode that sends the robot on an abrupt departure from the direction of any detected predator. Equal priority is given to all sensors.
3. **Prey sensing module** — an 8x2 ANN that takes input from the “infrared sniffing” mode (preliminary experiments investigated the effect of evolving this module).
4. **Wander module** — a programmed module that determined the default behaviour and speed of robot as well as sending it in random directions. This operated if none of the other modules were triggered.
5. **Light detection module** — an 8x2 ANN that takes input in the form of light intensity with two outputs to the motors to enable light following. This module is triggered by excessive background light on the IR sensors.
6. **Light tree module** — a 4x2 ANN that takes voltage change on the quadrants of the solar panels and steers the robot to the centre of the light source above it.

The predator behaviour is controlled by six modules:

1. **Object module** — an 8x2 ANN similar to that used in the prey but designed to take the shape of the predator into account.
2. **Predator sensing module** — an 8x2 ANN that takes input from the “infrared sniffing” mode that enables the predators to avoid each other at greater distances than in the IR distance-sensing mode.
3. **Wander module** — as in the prey, it sets default behaviour, speed, and direction of the robot.
4. **Prey sensing module** — an 8x2 ANN that takes input from the “infrared sniffing” mode and turns the robot in the direction of the sensed prey.
5. **Attentional module** — a 3x2 ANN that takes input from the front three IRs in “sniffing” mode. This has a sequential dependency on the prey-sensing module, and is only triggered when a prey has been detected and the predator has been orientated to face in its direction. Once activated, the attentional module entirely directs the behaviour of the robot while the prey is in its sights. If it loses the prey, it starts again. The attentional module was required because without it, the predator would keep swapping which prey it was chasing.
6. **Kill module** — essentially a programmed *reflex*. When a prey hits the front inset bump sensors, it triggers the lifting mechanism which lifts the robot into the air and initiates the movement of the copper-tipped spike into the brass plate of the prey. The spike is connected directly onto a circuit on the predator that drains battery from the prey. Lifting in the air was not necessary for the energy transfer but was added to heighten the drama of the museum show.

The application worked well for the clients at Magna. It was so popular with the media and the public that there were 7x30 minute shows per day. Thirty thousand people came to see the shows in the first week and 400,000 in the first six months. Each show started with an informal family-friendly lecture, and then the public watched a 20-minute free run of the robots. The number of robots varied and depended on the number of breakdowns. During very busy periods, three predators and eight prey were used so that there were two spare predators and two spare prey.

The number of catches per show ranged from five to 27. The average number was 11 in the 20 minute period. Interestingly, it was not the number of predators that made the difference, but rather which predator. Although they were all made to identical specifications, there were minor differences between them. For example, one of the predators, Gaak, became famous for its catches. It always caught more than any of the others even with a less good control algorithm.

On first consideration, the robots described in this study do not seem particularly swarm-like, in terms of their size and number. In other respects though, they can be seen to adhere to swarm-intelligent principles. They are biologically inspired (*Biological Inspiration* section), albeit in a generalised, not a specific sense, since they are not based on any particular animal or insect. They are individually-simple robots in that they are controlled in a reactive manner, using a combination of behavioural modules (*Individual Simplicity* section). The behavioural modules took the form of handcrafted weights in artificial neural networks (their values set on the basis of trial and error experimentation, and then evolved to a certain extent). They are clearly autonomous: not subject to remote or global control (*Collective Control* section). They are capable of only limited and local communication (*Communication* section); the infrared sniffing to detect predators and prey constitutes a form of robot interaction via sensing (*Communication* section). Their effective operation in the displays — avoiding obstacles and catching prey — provides a good illustration of the apparently complex behaviour that can result from simple reactive mechanisms. They do, however, represent an unusual use of robots, in that their application is one of educational entertainment in a museum. However, it is apparent that such an application can provide a useful source of funding for the further development of robotics.

CONCLUSION

In this chapter, we have considered what it means to take a swarm intelligence approach to collective robotics; we also looked at recent research and a case study. In this final section, we shall review the strengths and limitations of the approach, and its likely future directions.

Strengths and Limitations of the Approach

The main strengths of the swarm-based approach to robotics stem from its design philosophy, and reliance on the use of simple expendable robots that can accomplish a task through emergent cooperation. Swarm intelligence is a biologically-inspired notion, and rests on the concept of simple autonomous agents, with an emphasis on decentralised control and communication. Intelligent behaviour emerges from the interactions of such agents with each other, and with their environment. When the notion of swarm intelligence is realised in physical robots, it can be termed “swarm robotics”. The main characteristics of swarm robotics are:

1. that it is biological inspired (in a general, rather than a specific sense, Sharkey, 2003);
2. simplicity at the individual level (robots controlled by means of a system of behavioural or reactive modules that enable a fast response to the environment);

3. the robot collective is not centrally controlled (the robots are autonomous) and
4. the robots are capable of only local or indirect communication.

There are several strengths that can be derived from these characteristics. The principle motivation of swarm robotics is based on the idea that by emulating some of the characteristics of biological swarms, some of their emergent capabilities might be achieved. A swarm of autonomous redundant robots could be used to accomplish tasks in areas that are inaccessible or hazardous to humans. Simplicity and reactive control at the level of individual robots could mean that they would be able to respond rapidly and flexibly to an unknown and changing environment. Local control and communication means that individual robots could be added or removed from the collective without the risk of task failure or the need for recalibration. Similarly, the use of local rather than global control and communication means that there would be no problems with a communication bottleneck, or with the failure of a centralised processor.

However, applying a swarm intelligence approach to robot collectives also necessarily engenders some limitations. For a start, it implies the need to limit, or restrict, the abilities of the individual robots. The insistence on autonomy, simplicity, and local communication restricts the kind of task for which such robots could be used; complex tasks that require the detailed and centralised communication of changing instructions, for example, will not be appropriate. In addition, since part of the promise of swarm robotics is the notion of emergent behaviour, there is the question of finding ways of achieving the desired emergent behaviour. There are as yet no straightforward methods for designing the component behaviours that will result in the emergent behaviour required for the performance of user-designated tasks. Obtaining a solution for a particular task is likely to require a considerable amount of experimentation, to find the best robots, behaviours, and control systems to accomplish it in a particular environment. Possible methods (evolutionary algorithms, and neural network learning) were briefly outlined in the chapter, but represent only a starting point for solutions. Other limitations stem from the limited abilities of current robots. Clearly, there are many insect behaviours of which they are incapable, for instance, in terms of their mobility and sensors. Nonetheless, the promise is that there are some tasks that can be better solved through a swarm intelligent approach, than by methods involving centralised processing and control, and fewer more sophisticated robots.

Future research is likely to progress some of these areas towards more complete applications, although it will be interesting to see the extent to which a purist approach will be maintained, or whether a swarm-based approach will inevitably be augmented by direct control methods. One way of doing this would be to take a hybrid approach, and incorporate some more sophisticated robots into a collection. For example, it could prove useful to use a more sophisticated robot to shepherd, or transport, a swarm of simpler robots to a required location to complete a task. It is also likely that, in the future, attempts will be made to incorporate some more aspects of the organisation of insect societies into swarm robotics. Such efforts are beginning to be apparent in the explorations of the mechanisms of task allocation, and of caste divisions and dominance hierarchies reported earlier. Further research could also be legitimately used to extend the capabilities of individual robots. It could be used legitimately in the sense of still being biologically plausible, since although insect societies are clearly organised on

decentralised principles, it is quite likely that individual insects are capable of more than purely reactive behaviour. There is evidence that they can form some kinds of representation, and can communicate with each other to a greater extent than has yet been modelled in robotics; bees, for example, can recognise landmarks, and can perform dances that let other members of the hive know where to find good sources of food outside the nest. The possibility of incorporating a more detailed knowledge of insect behaviours into swarm robotics is one that may be realised in the future.

REFERENCES

- Agassounon, W., Martinoli, A., & Easton, K. (2004). Macroscopic modeling of aggregation experiments using embodied agents in teams of constant and time-varying sizes. In M. Dorigo & E. Sahen (Eds.), *Autonomous Robots (special issue on Swarm Robotics)*, 17(2-3), 163-191.
- Arkin, R. C. (1990). Integrating behavioural, perceptual and world knowledge in reactive navigation. *Robotics and Autonomous Systems*, 6, 105-122.
- Arkin, R. C. (1998). *Behaviour-based robotics*. Cambridge, MA: MIT Press.
- Arkin, R. C. (2003). Reactive robotic systems. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (2nd ed.). Cambridge, MA: Bradford Books/MIT Press.
- Balch, T. (2002). Measuring robot group diversity. In T. Balch & L. Parker (Eds.), *Robot teams: From diversity to polymorphism*. Natick, MA: A. K. Peters.
- Baldassarre, G., Nolfi, S., & Parisi, D. (2002). Evolving mobile robots able to display collective behaviours. *Proceedings of the International Workshop on Self-Organisation and Evolution of Social Behaviour*, Monte Verita, Switzerland, September.
- Beckers, R., Holland, O. E., & Deneubourg, J. L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In R. Brooks & P. Maes (Eds.), *Artificial life IV*. Cambridge, MA: MIT Press.
- Beni, G., & Wang, J. (1989). Swarm intelligence. *Proceedings of the 7th Annual Meeting of the Robotics Society of Japan*, Tokyo, Japan (pp. 425-428).
- Birk, A., & Belpaeme, T. (1998). A multi-gent system based on heterogeneous robots. In A. Drogoul, M. Tambe, & T. Fukuda (Eds.), *Proceedings of the 1st International Collective Robotics Workshop — CRW 98, Lecture Notes in Artificial Intelligence* (Vol. 145, pp. 13-24). Berlin, Germany: Springer-Verlag.
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. Oxford, UK: Oxford University Press.
- Bonabeau, E., & Theraulaz, G. (2000). Swarm smarts. *Scientific American*, 282(3), 72-9.
- Bourke, A. F. G., & Franks, N. R. (1995). *Social evolution in ants*. Princeton, NJ: Princeton University Press.
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
- Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 2, 14-23.
- Brooks, R. (1991). New approaches to robotics. *Science*, 253, 13 September, 1227-1232.
- Brooks, R. (1999). *Cambrian intelligence: The early history of the new AI*. Cambridge, MA: Bradford Books/MIT Press.

- Burgard, W., Moors, M., Fox, D., Simmons, R., & Thrun, S. (2000). Collaborative multi-robot exploration. *Proceedings of the IEEE International Conference on Robotics and Automation*, San Francisco, April 24-28, Los Alamitos, CA.
- Cao, Y. U., Fukunaga, A. S., & Kahng, A. B. (1997). Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4, 7-27.
- Desai, J., Ostrowski, J., & Kumar, V. (1998). Controlling formations of multiple mobile robots. *Proceedings of the IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May 16-20 (pp. 2864-2869).
- Dorigo, M., & Colombetti, M. (1998). *Robot shaping: An experiment in behavior engineering*. Cambridge, MA: Bradford Books/MIT Press.
- Dorigo, M., & Gambardella, L. M. (1997). Ant colony system: A cooperative learning approach to the travelling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1), 53-66.
- Dudek, G., Jenkin, M., & Milios, E. (2002). A taxonomy of multi-robot systems. In T. Balch & L. Parker (Eds.), *Robot teams: From diversity to polymorphism*. Natick, MA: A. K. Peters.
- Dudek, G., Jenkin, M., Milios, E., & Wilkes, D. (1996). A taxonomy for multi-agent robotics. *Autonomous Robots*, 3(4), 375-397.
- Gerecke, U., Sharkey, N. E., & Sharkey, A. J. C. (2000). Reliable robot localization with an ensemble. *Proceedings of the 2nd International Symposium on Robotics and Automation (ISRA-2000)*, Monterrey, Mexico, November (pp. 515-520).
- Goldberg, D., & Mataric, M. (2002). Design and evaluation of robots behavior-based controllers. In T. Balch & L. Parker (Eds.), *Robot teams: From diversity to polymorphism*. Natick, MA: A. K. Peters.
- Grabowski, R., Navarro-Serment, L. E., Panedis, C. J. J., & Khosla, P. K. (2002). Heterogeneous teams of modular robots. In T. Balch & L. Parker (Eds.), *Robot teams: From diversity to polymorphism*. Natick, MA: A. K. Peters.
- Grassé, P. P. (1959). La reconstruction du nid et les coordinations inter-individuelles chez *bellicositermes natalensis* et *cubitermes* sp la théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeur. *Insectes Sociaux*, 6, 41-80.
- Holland, O., & Melhuish, C. (1999). Stigmergy, self-organisation, and sorting in collective robotics. *Artificial Life*, 5, 173-202.
- Ijspeert, A. J., Martinoli, A., Billard, A., & Gambardella, L. M. (2001). Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots*, 11, 149-171.
- Jung, D., & Zelinsky, A. (1999). An architecture for distributed cooperative planning in a behaviour-based multi-robot system. *Robotics and Autonomous Systems*, 26, 149-174.
- Krieger, M. J. B., & Billeter, J. B. (2000). The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, 30, 65-84.
- Kube, C. R., & Bonabeau, E. (2000). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30, 85-101.
- Kube, C., & Zhang, H. (1996). The use of perceptual cues in multi-robot box-pushing. *Proceedings of the IEEE International Conference on Robotics and Automation*, Minneapolis, MN, 22-28 April (pp. 2085-2090).

- Maes, P. (1989). The dynamics of action selection. *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, MI, August (pp. 991-997).
- Martinoli, A. (1999). *Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed collective strategies*. PhD thesis NR 2069, DI-EPFL, Lausanne, Switzerland.
- Mataric, M. (1992). Minimizing complexity in controlling a collection of mobile robots. *Proceedings of the IEEE International Conference on Robotics and Automation*, Nice, France, May (pp. 830-835).
- Mataric, M. (1997a). Reinforcement learning in the multi-robot domain. *Autonomous Robotics*, 4(1), 73-83.
- Mataric, M. (1997b). Behaviour-based control: Examples from navigation, learning, and group behaviour. *Journal of Experimental and Theoretical Artificial Intelligence (Special Issue on Software Architectures for Physical Agents)*, 9(2-3), 323-336.
- Murphy, R. (2002). Marsupial robots. In T. Balch & L. Parker (Eds.), *Robot teams: From diversity to polymorphism*. Natick, MA: A. K. Peters.
- Nolfi, S. (1997). Using emergent modularity to develop control systems for mobile robots. *Adaptive Behavior (Special issue on Environment Structure and Behavior)*, 3-4, 343-364.
- Nolfi, S. (2002). Evolving robots able to self-localize in the environment: The importance of viewing cognition as the result of processes occurring at different time-scales. *Connection Science*, 14(3), 231-244.
- Nolfi, S. (2003). Evolution and learning in neural networks. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (2nd ed.) (pp. 415-418). Cambridge, MA: Bradford Book/MIT Press.
- Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organising machines*. Cambridge, MA: Bradford Books/MIT Press.
- Nolfi, S., & Parisi, D. (1997). Learning to adapt to changing environments in evolving neural networks. *Adaptive Behaviour*, 1, 75-98.
- Pagello, E., D'Angelo, A., Montesello, F., Garelli, F., & Ferrari, C. (1999). Cooperative behaviours in multi-robot systems through implicit communication. *Robotics and Autonomous Systems*, 29, 65-77.
- Parker, L. E. (1994). ALLIANCE: Architecture for fault-tolerant cooperative control of heterogeneous mobile robots. *Proceedings of the IEEE/RSJ Intelligent Robotics and Systems Conference (IROS)*, Munich, Germany, September (pp. 776-783).
- Parker, L. E. (1998). ALLIANCE: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2), 220-240.
- Payton, D. W., Rosenblatt, J. K., & Keirsey, D. M. (1990). Plan guided reaction. *IEEE Transactions on Systems Man and Cybernetics*, 20(6), 1370-1382.
- Raignier, A., & van Boven, J. (1955). Étude taxonomique, biologique et biométrique des *Dorylus* du sou-genre *Anomma* (Hymenoptera Formicidae). *Annales de Musée Royal du Congo Belge*, n.s. 4, Sciences Zoologiques, 2, 1-359.
- Rekleitis, I., Dudek, C., & Milos, E. (1997) Multi-robot exploration of an unknown environment, efficiently reducing the odometry error. *Proceedings of the International Joint Conference on Artificial Intelligence — IJCAI* (pp. 1340-1345). San Francisco: Morgan Kaufmann.

- Schoonderwoerd, R., Holland, O.E., Bruten, J., & Rothkrantz, L. (1996). *Ant-based load balancing in telecommunications networks* (Tech. Rep. No. HPL-96-76, May 21), Hewlett-Packard Labs.
- Sharkey, A. (Ed.). (1999). *Combining artificial neural nets: Ensemble and modular multinet systems*. Berlin, Germany: Springer-Verlag.
- Sharkey, N. (1998). Learning from innate behaviours: A quantitative evaluation of neural network controllers. *Autonomous Robotics*, 5, 317-334.
- Sharkey, N. (2003). Biologically inspired robotics. In M. A. Arbib (Ed.), *The handbook of brain theory and neural networks* (2nd Ed.) (pp. 160-164). Cambridge, MA: Bradford Book/MIT Press.
- Spires, S. V., & Goldsmith, S. Y. (1998). Exhaustive geographic search with mobile robots along space-filling curves. In A. Drogoul, M. Tambe, & Fukuda, T. (Eds.), *Proceedings of the 1st International Collective Robotics Workshop — CRW98, Lecture Notes in Artificial Intelligence* (Vol. 145, pp. 1-12). Berlin, Germany: Springer-Verlag.
- Steels, L. (1996). Emergent adaptive lexicons. In P. Maes, J. –A. Meyer, & S. Wilson (Eds.), *From Animals to Animals 4: Proceedings of the 4th International Conference on Simulation of Adaptive Behaviour, Complex Adaptive Systems* (pp. 562-567). Cambridge, MA: MIT Press.
- Stuart, A. M. (1967). Alarm, defence, and construction behavior relationships in termites (Isoptera). *Science*, 156, 1123-1125.
- Walter, W. G. (1953). *The living brain*. New York: Norton.
- Watkins, C. J. C. H. (1989). *Learning with delayed rewards*. PhD dissertation, King's College, University of Cambridge, UK.
- Watson, R. A., Ficici, S. G., & Pollack, J. B. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39, 1-18.
- Wilson, E. O. (1971). *The insect societies*. Cambridge, MA/London: The Belknap Press of Harvard University Press.
- Yamauchi, B. (1999). Decentralized coordination for multirobot exploration. *Robotics and Autonomous Systems*, 29, 111-118.

Chapter VII

Self-Organising Impact Sensing Networks in Robust Aerospace Vehicles

Mikhail Prokopenko, Geoff Poulton

Don Price, Peter Wang

Philip Valencia, Nigel Hoschke

Tony Farmer, Mark Hedley

Chris Lewis, and Andrew Scott

CSIRO Information and Communication Technology Centre and

CSIRO Industrial Physics, Australia

ABSTRACT

An approach to the structural health management (SHM) of future aerospace vehicles is presented. Such systems will need to operate robustly and intelligently in very adverse environments, and be capable of self-monitoring (and ultimately, self-repair). Networks of embedded sensors, active elements, and intelligence have been selected to form a prototypical “smart skin” for the aerospace structure, and a methodology based on multi-agent networks developed for the system to implement aspects of SHM by processes of self-organisation. Problems are broken down with the aid of a “response matrix” into one of three different scenarios: critical, sub-critical, and minor

damage. From these scenarios, three components are selected, these being: (a) the formation of “impact boundaries” around damage sites, (b) self-assembling “impact networks”, and (c) shape replication. A genetic algorithm exploiting phase transitions in systems dynamics has been developed to evolve localised algorithms for impact boundary formation, addressing component (a). An ant colony optimisation (ACO) algorithm, extended by way of an adaptive dead reckoning scheme (ADRS) and which incorporates a “pause” heuristic, has been developed to address (b). Both impact boundary formation and ACO-ADRS algorithms have been successfully implemented on a “concept demonstrator”, while shape replication algorithms addressing component (c) have been successfully simulated.

INTRODUCTION

Structural health management (SHM) is expected to play a critical role in the development and exploitation of future aerospace systems, operating in harsh working environments and responding to various forms of damage and possible manufacturing and/or assembly process variations. SHM is a new approach to monitoring and maintaining the integrity and performance of structures as they age and/or sustain damage. It differs from the traditional approaches of periodic inspection and out-of-service maintenance by aiming for continuous monitoring, diagnosis, and prognosis of the structure while it is in service, damage remediation and, ultimately, self-repair. This requires the use of networked sensors and active elements embedded in the structure, and an intelligent system capable of processing and reducing the vast quantities of data that will be generated, to provide information about the present and future states of the structure, and to make remediation and repair decisions.

This chapter outlines an approach being taken to the development of next-generation SHM systems, and the development of a flexible hardware test-bed for evaluating and demonstrating the principles of the approach. This introductory section will outline the general requirements of an SHM system, provide an overview and relevant details of the hardware test-bed, and introduce our approach to the systems-level issues that must be solved.

Structural health management systems will eventually be implemented in a wide range of structures, such as transport vehicles and systems, buildings and infrastructure, and networks. Much of the current research effort is aimed at the highvalue, safety-critical area of aerospace vehicles. CSIRO is working with NASA (Abbott, Doyle, Dunlop, Farmer, Hedley, Herrmann et al., 2002; Abbott, Ables, Batten, Carpenter, Collings, Doyle et al., & Winter, 2003; Batten, Dunlop, Edwards, Farmer, Gaffney, Hedley et al., 2004; Hedley, Hoschke, Johnson, Lewis, Murdoch et al., & Farmer, 2004; Price, Scott, Edwards, Batten, Farmer, Hedley et al., 2003; Prokopenko, Wang, Price, Valencia, Foreman, Farmer, 2005a) and other key industry players to develop and test concepts and technologies for next-generation SHM systems in aerospace vehicles. While many of the principles of SHM systems described in this chapter are quite general, aerospace vehicles will be used throughout as example structures.

General Requirements of a Structural Health Management System

The key requirements of an advanced health monitoring system are that it should be able to detect damaging events, characterize the nature, extent, and seriousness of the damage, and respond intelligently in whatever timescale is required, either to mitigate the effects of the damage or to effect its repair. Strictly speaking, a pure monitoring system is expected only to report damage rather than to formulate a response, but it is preferable that the ultimate objective of responding to damage be borne in mind from the outset.

The statement of key requirements serves to sub-divide the problem as follows:

1. **Detection of damaging events**, which requires some knowledge of the environment in which the vehicle will be operating, the threats that it will face, and the development of sensors as well as a strategy for using them to detect damage events well within the time required for the system to respond. For events that require a rapid response, the best solution will often involve the use of passive, embedded sensors.
2. **Evaluation of the extent and severity of the damage**. This may or may not be a separate process from event detection. It may use different sensors, or the same sensors may be used in a different way. It is more likely to employ active sensors, which may be embedded in the structure or could be mobile and autonomous.
3. **Diagnosis of the damage**, which includes identification of the nature of the damage (for example, is it due to corrosion, fatigue, impact, and so on?) and its cause. An intelligent system should be able to utilize data from a vast array of sensors to deduce information about the events that have occurred and the resulting damage, on a whole-of-vehicle basis. Knowledge of the cause of damage may enable actions to be taken to reduce the rate of damage progression. Diagnosis also requires an assessment of the effect of the damage on the performance capability and integrity of the structure.
4. **Prognosis for the structure** requires prediction of the future progression of the damage and assessment of the effect of the forecast damage on structural performance. It requires an estimate of the future operating conditions of the structure.
5. **Formulation of the response: intelligent decision-making**. The nature of the response will depend on a number of factors such as the range of possible response mechanisms, the diagnosis of the damage (steps 2 and 3 above), the available response time as deduced from the diagnosis and prognosis (step 4), and so on. A response may consist of a sequence of actions. Major damage may demand an immediate emergency response, such as the rapid isolation of a whole section of the vehicle, followed by a more considered damage evaluation and repair strategy.
6. **Execution and monitoring of the response**. In addition to repair, a holistic response may involve changes to the flight or operational characteristics of the vehicle, either to mitigate the effects of the damage or to assist in the avoidance of further damage. The effectiveness of the response will require monitoring.

The first and second of these points are what is generally referred to as structural health monitoring. It is currently carried out in a very limited way in specific regions of selected structures (for instance, some aircraft, some items of large infrastructure),

generally using a small number of sensors connected to a data logger or computer. Ultimately, large numbers of sensors will be required to detect and evaluate a wide range of possible damage types within a large and complex structure.

NASA's vision of self-monitoring robust aerospace vehicles includes both local and global SHM systems (Generazio, 1996). The local actions are anticipated to autonomously identify, evaluate, and trigger an appropriate response, including repair, for a wide range of damage or defect conditions in aerospace materials and structures, using distributed micro-sized sensors, multiple miniature robotic agents, micro-sized repair tools, and self-healing smart structures. In parallel, global actions should enable dynamic evaluation of structural integrity across large and remote areas. This dual architecture, in turn, entails the need for dynamic and decentralised algorithms used in all the key requirements enumerated above.

An additional key requirement of an autonomous SHM system is robustness. The system must be able to operate effectively in the presence of damage to the structure and/or failure of system components: its performance must degrade "gracefully" rather than catastrophically when damage occurs. Scalability, reliability, and performance verification are also needed.

Also of great importance to any SHM system is the provision of an efficient and robust communications system. Unless local actions are sufficient, the key requirements mentioned above will rely on communication from a damage site to another part of the vehicle, for example, to initiate secondary inspections, repair, or in extreme cases, appropriate emergency action. Such communications will most likely be hierarchical and flexible, since the site to which damage is reported will vary with time, as well as the damage location and severity. Robustness must also be a feature, with continuing communications ensured even in severe damage situations.

In order to address these requirements, we have chosen to apply a multi-agent system (MAS) approach to the architecture, and seek to develop design methodologies that will enable the desired responses of the system (the remedial actions) to emerge as self-organised behaviours of the communicating system of sensing and acting agents. The particular MAS structure which is the focus of this chapter is a group of contiguous agents, locally connected and forming the surface of a three-dimensional object. Each agent has sensing and computational capabilities, and can communicate only with its immediate neighbours. Thus all communications, local, regional, and global need to occur through these agent-to-agent links. Although such constraints impede the flow of information, there is a significant potential redundancy which can aid robustness. Various types of communications will be needed, ranging from local cell-to-cell handshaking to check status, to emergency global communications in case of severe damage, which must be carried out as rapidly as possible whenever needed.

Much of this research has been undertaken as part of the CSIRO-NASA Ageless Aerospace Vehicle (AAV) project, which also includes an experimental test-bed and concept demonstrator (CD) system, whose aim is to detect, locate, and evaluate impacts by fast particles. A software simulation package has also been developed. The purpose of these two tools is to provide versatile research platforms for investigations of sensing, data processing, communications, and intelligence issues, and for demonstrating solutions for some of these issues. The architecture of the system is highly modular, being composed of "cells" that constitute the outer skin of the vehicle. Each cell consists of

a small region of the vehicle skin, a number of sensors attached to the skin, a processing unit, and communication ports. Each of these cells is an agent in the multi-agent system architecture. This system will be described in more detail in the *Ageless Aerospace Vehicle Project* section.

The Approach to Intelligent SHM System Development

The approach adopted here to the development of system intelligence is based on a multi-agent system (Ferber, 1999) in which the desired responses emerge by self-organisation. What is meant by self-organisation? The following definition, in the context of pattern formation in biological systems, was given by Camazine, Deneubourg, Franks, Sneyd, Theraulaz, and Bonabeau, (2001):

Self-organization is a process in which pattern at the global level of a system emerges solely from numerous interactions among the components of the system. Moreover, the rules specifying interactions among the system's components are executed using only local information, without reference to the global pattern.

This definition captures two important aspects of self-organisation. Firstly, the global behaviour of the system of many interacting components (agents) is a result only of the interactions between the agents, and secondly, that the agents have only local information, and do not have knowledge of the global state of the system. Typically, this emergent behaviour at the system level is not easily predictable from local agents' rules and interactions.

Self-organisation occurs in both biological and non-biological systems. In non-biological systems, self-organisation is produced by a flow of energy into the system that pushes it beyond equilibrium: the winds that produce characteristic ripples in sand, the temperature gradients that produce Bénard convection cells in a viscous fluid, the thermodynamic forces that lead to crystal growth and characteristic molecular conformations are all examples of these external energy inputs. However, the nature of the emergent behaviour depends critically on the interactions between the low-level components of the systems — the grains of sand, the molecules in the fluid, the atoms in the crystals and molecules. These interactions are determined by the laws of nature and are immutable.

In biological systems, on the other hand, the interactions between components of a system may change over generations as a result of evolution. There are selection pressures based on adaptation to the environment and survival. These selection pressures lead to emergent behaviour that is desirable for the survival of the system in the environment in which it has evolved, but which may be undesirable in other environments. Similarly, when using evolutionary methods for the design of complex SHM systems that employ self-organised responses to damage, there is a need to identify appropriate selection pressures. These, through their contribution to an evolutionary fitness function, will constrain the agent interactions to produce desirable emergent responses. Such selection pressures will be further discussed later in this chapter.

Current approaches developed for complex systems, and in particular, multi-agent networks, either solve individual problems using evolutionary algorithms, or restrict the

solution space so that emergent behaviour is impossible. Both these approaches are inadequate, the first because of high computational needs and the loss of an intuitive feel for the results, and the second because it is likely to over-constrain the range of possible solutions: it is noted that biological systems make extensive use of the rich solution space provided by the complexity of natural systems. In order to emulate this capability, a general design methodology, retaining the essential complex behaviour of multi-agent systems, is needed. Design in this context means the ability to specify the local agent properties so that they interact to produce a desired global result.

In this chapter, we describe an initial hybrid top-down/bottom-up (TDBU) attempt at subdividing a set of high-level goals into intermediate hierarchical objectives, and exploring the solution space at each intermediate level of the hierarchy. In particular, we explicitly define the main functional SHM sub-tasks, working downwards from the top-level design goals. The next stage is, for each sub-task, to design localised algorithms working from the bottom up and using an iterative process including the following steps:

1. forward simulation leading to emergent behaviour for a task-specific class of localised algorithms;
2. quantitative measurement of desirable qualities shown by the emergent patterns (for example, spatiotemporal stability, connectivity, and so on); and
3. evolutionary modelling of the algorithms, with the metrics obtained at step (b) contributing to the fitness functions.

While the eventual optimal solution to the overall SHM problem may not involve sequential steps through the sub-tasks listed above, our initial approach is to divide the problem along the lines indicated. Thus we will first aim to develop procedures to characterise damage (in terms of its nature, location, extent, and severity), then form a diagnosis, then a prognosis, and finally make decisions and take appropriate actions.

A diagnosis, or the confidence in a diagnosis, may change with time, as the development of damage is monitored and more information becomes available. One of the major benefits of SHM is the ability it provides to detect damage at an early stage and to monitor its development, leading to improved diagnostic capabilities and, ultimately, more efficient repair strategies. Similarly, a prognosis, which depends on prediction of the future progression of the damage, can be modified with time as the damage develops.

The Response Matrix Approach to Comparing Response Characteristics

It is clear that an intelligent system in a safety-critical environment must be able to respond very differently in different circumstances. In the event of sudden critical damage, such as a major impact, the most important characteristic of the response may be speed. Some undesirable side effects may be a tolerable trade-off for a rapid and effective emergency response. On the other hand, an acceptable response to slowly developing non-critical damage, such as highcycle fatigue or corrosion, must be more deliberative and targeted, and response speed is unlikely to be a relevant consideration. In order to provide a basis for comparison of response types, and to guide thinking about the processes by which responses are produced, the following simple response matrix method has been developed.

The response matrix seeks to classify a response on the basis of its spatial extent and the degree of deliberation required to form the response. The spatial extent is defined in terms of system cells, where a cell is the smallest intelligent unit of the system. Examples of cells in the AAV test-bed are described in the *Ageless Aerospace Vehicle* section. A response is categorized as “local” if only a single cell is involved, as “neighbourhood” if only a small group of neighbouring cells is involved, or as “global” if a larger region, such as a complete sub-structure or even the whole structure is involved in the response. The nature of a response is considered to be “reactive” if it is made rapidly, using only the initially-sensed data, and with effectively no feedback that could be classed as deliberation. It is said to be “strongly deliberative” if there are long feedback loops involved in obtaining additional sensed data, and making a response that would be classified as being highly intelligent. A “moderately deliberative” response would involve some deliberation by the system, but with shorter feedback loops than required for a strongly deliberative response. Some examples of these responses will be outlined below to clarify these definitions; but first, three levels of damage will be defined.

Three levels of damage will be referred to throughout this chapter. The first is *critical damage*, which is sufficiently severe to threaten the integrity of the structure, and possibly the survival of the vehicle. Critical damage will generally occur suddenly, or presumably its precursors would have been detected and corrected. It will require an emergency response that must be rapid and effective, even if subsequent, more thorough diagnosis shows it to have been an over-reaction. The second level is *sub-critical damage*, which, although severe enough to require an immediate response, is not sufficiently threatening to the vehicle’s survival to require an emergency response. Thirdly, there is *non-critical, or minor damage*, which does not necessarily require an immediate response, but which must be monitored to track its progression with time (as with, for example, corrosion or fatigue damage), or its possible interaction with other damage mechanisms.

In terms of these levels of damage, a reactive response will generally be invoked only by critical damage, or by an indication of the likelihood of critical damage: it will generally be preferable to react to the likelihood of critical damage than not. A reactive response is pre-programmed (such as an emergency evacuation from a building), and will be followed by a more deliberative evaluation and diagnosis. It may include physical and/or electronic isolation of a cell, neighbourhood, or sub-structure, and the initiation of autonomic and fast temporary repairs. The response to indications of sub-critical or non-critical damage will be to evaluate the severity of the damage, by monitoring the outputs of sensors other than those that indicated the damage, or by initiating active damage evaluation (ADE) using either embedded or mobile sensors. The ADE and subsequent remedial actions may be moderately deliberative or strongly deliberative, depending on the amount of information and prior knowledge which is required for a diagnosis. A moderately deliberative response might consist of a rapid diagnosis from a single set of ADE data, followed by an immediate remedial response. On the other hand, an accurate diagnosis of non-critical damage might require the damage progression to be monitored for some time, or it may require several sets of ADE data and comparison with a physical damage model, and this would be considered a strongly deliberative response.

Examples of the ways in which chains of responses to different types and levels of damage can be classified using the response matrix approach are shown in Table 1. Two

forms of potentially critical damage, a fast particle (perhaps a micro-meteoroid) impact on a single cell, and large body impact or an explosion that causes severe damage to a whole sub-structure, as well as one of non-critical damage to one or more cells, are considered. In the cases of critical damage, the initial reactive response that is required

Table 1. The response matrix approach for classification of system responses, with examples for three types of damage, as outlined in the text

Degree of Deliberation	Reactive	<u>Major damage to single cell by fast particle impact.</u> <ul style="list-style-type: none"> – Ignore messages from damaged cells. – Report damage to neighbourhood action point(s). – Initiate automatic temporary seal. Follow by Moderately Deliberative, Neighbourhood response.		<u>Major damage to sub-structure by large body impact or by explosion or fire.</u> <ul style="list-style-type: none"> – Initiate emergency comms. to global action site. – Physically and electronically isolate sub-structure. Follow by Moderately Deliberative, Global response
	Moderately Deliberative		<ul style="list-style-type: none"> – ADE by neighbouring cells to assess damage to cell. – If damage not critical, and/or if local diagnostics are favorable, re-enable cell. – If damage critical, initiate cell replacement. 	<ul style="list-style-type: none"> – ADE by remote sensors, mobile sensors or robotic swarm. – Identify damaged region. – Re-enable cells that are not critically damaged. Follow by Strongly Deliberative, Global response
	Strongly Deliberative	<u>Corrosion or degradation damage to single cell.</u> <ul style="list-style-type: none"> – Monitor damage and environmental effectors as damage progresses. 	<ul style="list-style-type: none"> – Set up damage network in neighbourhood. – Compare with damage models or prior knowledge to develop diagnosis and prognosis. – Take remedial action when indicated by prognosis. 	<ul style="list-style-type: none"> – Diagnosis and prognosis of damaged region. – Respond as indicated by prognosis: repair or replace damaged cells if necessary, monitor state of others.
		Local	Neighbourhood	Global
Spatial Extent				

to ensure survival is followed by a more deliberative response to obtain more specific information about the damage, to produce more appropriate long-term remediation of the damage, and to enable the system to learn to deal better with similar events in the future. Similar sequences of reactive and deliberative response to danger (“panic” response) and damage can be recognized in animals, including humans.

Top-Down/Bottom-Up Design (TDBU) and the Response Matrix

One way of viewing the response matrix is as the top-down part of the TDBU approach to design which was outlined in the *Approach to Intelligent SHM System Development* section. For a range of damage scenarios and the desired system response to each, the response matrix infers the large-scale components necessary for the appropriate response to occur. The components, whilst not unique, are chosen as high-level and as broad in spatial extent as possible, in line with the minimal hierarchical decomposition of the problem which is the intention of the TDBU approach.

It is envisaged that most if not all of the components would be implemented by self-organisation within the multi-agent structure. If the decomposition is too broad, then there may be difficulty in achieving such self-organising solutions, while if it is too prescriptive, then the result may be an unsatisfactory system outcome. Of course, the ideal would be to achieve complete, self-organised responses to all likely damage scenarios without having to decompose the problem. The possibility of this is unlikely, at least in the near future, because of the complexity of multi-agent systems, so a minimal hierarchical decomposition is a good compromise.

The most important components are damage detection, local assessment of damage, higher-level assessment (diagnosis and prognosis), and response (actions). Some of these will be treated in detail in later sections. However, there is one critical component which cannot be left out of the equation. This is communications, which is a necessary part of all damage scenarios, and at least as complex and difficult to handle properly as the others mentioned above.

Communications

Agents in a multi-agent system communicate either directly or through the environment (stigmergy) to form a network which usually exhibits complex behaviour (Holland & Melhuish, 1999). If the agents are fixed in space, as on the skin of an aerospace vehicle, then direct inter-agent communications forms the basis of the network. As will be seen, the particular agent networks of interest here generally only support communications between adjacent neighbours. Although this appears at first glance to be a restriction, it is also the main source of network robustness because of the large redundancy provided by the network.

There are many different communications tasks which the network has to be able to handle, ranging from simple status queries and responses for adjacent neighbours to the reporting of a critical damage situation to a remote site from which appropriate action can be initiated. Although communications tasks vary with the damage scenario, they all share the need to transfer information robustly and efficiently from one part of the network to another, in an environment where both transmit and receive sites may not know the other’s location (which may change with time anyway). In addition, the

environment itself is time-variable, particularly in times of significant damage. Secure communications in such an environment is a general task of significant difficulty which, however, needs to be solved since the survival of the vehicle may depend on it.

In a multi-agent system it would seem natural to seek self-organising solutions to this problem, and some progress has been made in this area, including one application presented in the *Impact Boundaries* section.

BACKGROUND AND RELATED WORK

Self-organisation is typically defined as the evolution of a non-equilibrium system into an organised form in the absence of external pressures. Over the last years, a number of examples employing self-organisation have been suggested in the broad context of biological and bio-inspired multi-agent systems: the formation of diverse spatial structures by groups of ants (Deneubourg & Goss, 1989), the growth and morphogenesis of networks of galleries in the ant *Messor sancta* (Buhl, Deneubourg, & Theraulaz, 2002); a propulsive motion of locally-connected mobile automata networks, dynamically organising into simple spatial structures while evolving toward task-specific topologies (Wessnitzer, Adamatzky, & Melhuish, 2001); a pattern formation of self-assembling modular robotic units, with the emergent chaining behaviour being analogous to the process of polymerisation, and the emergent clustering behaviour being similar to the autocatalytic process used by pheromone-depositing bark beetle larvae (Trianni, Labella, Gross, Sahin, Dorigo, & Deneubourg, 2002); fault-tolerant circuit synthesis on a self-configurable hardware platform provided by the Cell Matrix approach (Durbeck & Macias, 2002); a self-assembly of network-like structures connecting a set of nodes without using pre-existing positional information or long-range attraction of the nodes, using Brownian agents producing different local (chemical) information, responding to it in a non-linear manner (Schweitzer & Tilch, 2002).

Traditional multi-component systems do not exhibit self-organisation; instead, they rely on fixed multiple links among the components in order to efficiently control the system, having fairly predictable and often pre-optimised properties, at the expense of being less scalable and less robust. In the SHM context, condition-based maintenance (CBM), a process where the condition of equipment is automatically monitored for early signs of impending failure, followed by diagnostics and prognostics, has become popular for multi-component systems due to its cost and reliability advantages over traditional *scheduled* maintenance programs. However, according to a NASA Jet Propulsion Laboratory (JPL) report on Prognostics Methodology for Complex Systems (Gulati & Mackey, 2003), CBM is frequently difficult to apply to complex systems exhibiting emergent behaviour and facing highly stochastic environmental effects. A scalable solution capable of providing a substantial look-ahead capability is required. The JPL solution involves an automatic method to schedule maintenance and repair, based on a computational structure called the informed maintenance grid, and targeting the two fundamental problems in autonomic logistics: (1) unambiguous detection of deterioration or impending loss of function, and (2) determination of the time remaining to perform maintenance or other corrective action based upon information from the system (Gulati & Mackey, 2003). The solution based on the JPL work does not account for self-organisation and is not directly applicable to distributed multi-agent networks.

A recent paper by Prosser, Allison, Woodard, Wincheski, Cooper, Price, Hedley, Prokopenko, Scott, Tessler, and Spangler, (2004) has given an overview of NASA research and development related to SHM systems, and has discussed the requirements for SHM systems architectures. Characteristics such as scalability, flexibility and robustness were identified as being important requirements. Biological systems, including those referred to above, provide many examples of these characteristics in self-organising multi-agent systems. Indeed, it has been asserted that biological complexity and self-organisation have evolved to provide these characteristics. For example, Klyubin, Polani and Nehaniv (2004) indicated that evolution of the perception-action loop in nature aims at improving the acquisition of information from the environment and is intimately related to selection pressures towards adaptability and robustness — their work demonstrated that maximisation of information transfer can give rise to intricate behaviour, induce a necessary structure in the system, and ultimately be responsible for adaptively reshaping the system. In order to investigate the practical implementation of biologically-inspired concepts to structural health management systems, an experimental multi-agent test-bed has been developed. This will be described in the next section.

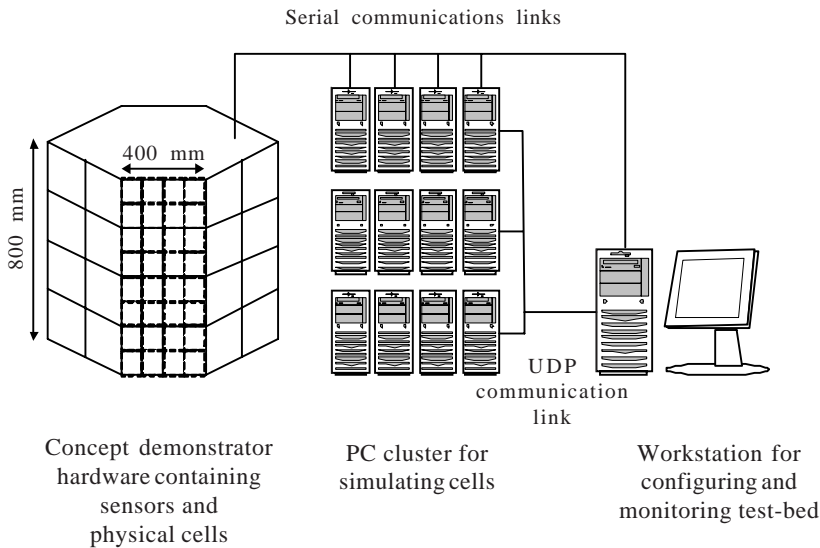
THE AGELESS AEROSPACE VEHICLE PROJECT

Introduction

The CSIRO-NASA Ageless Aerospace Vehicle (AAV) project has developed and examined concepts for *self-organising* sensing and communication networks (Abbott et al., 2002; Abbott et al., 2003; Price et al., 2003; Batten et al., 2004; Hedley et al., 2004; Prokopenko et al., 2005a). These concepts are being developed, implemented, and tested in an experimental test-bed and concept demonstrator: a hardware multi-cellular sensing and communication network whose aim is to detect and react to impacts by high-velocity projectiles that, for a vehicle in space, might be micro-meteoroids or space debris. High-velocity impacts are simulated in the laboratory using short laser pulses and/or steel spheres fired using a light-gas gun.

The test-bed has been built as a tool for research into sensor design, sensing strategies, communication protocols, and distributed processing using self-organising multi-agent systems. It has been designed to be modular and highly flexible: By replacing the sensors and their associated interface and data acquisition electronics, the system can be readily reconfigured for other applications.

Figures 1 and 2 contain a schematic overview of the system and photographs of its physical implementation, respectively. The physical structure is a hexagonal prism formed from a modular aluminium frame covered by 220 mm x 200 mm, 1-mm thick aluminium panels that form the outer skin of the structure. Each such panel contains four “cells”, and each of the six sides of the prism contains eight of these panels. The skin therefore consists of 48 aluminium panels and 192 cells. Cells are the fundamental building blocks of the system: they are the electronic modules containing the sensing, processing, and communication electronics. Each cell occupies an area of ~100 mm x 100 mm of the skin, mounted on the inside of which are four piezo-electric polymer (PVDF)

Figure 1. Architecture of the test-bed

sensors in a 60 mm square array, to detect the acoustic waves that propagate through the skin as a result of an impact.

Each cell contains two electronic modules (Figure 2), one of which acquires data from the sensors, while the other runs the agent software and controls the communications with its neighbouring cells. Importantly, a cell communicates only with four immediate neighbours. The test-bed does not employ centralised controllers or communication routers.

Also shown in Figure 1 are a PC cluster and a workstation. The cluster is used to simulate a larger network of cells, and is used for research into the emergent behaviour of multi-agent algorithms in very large networks. The workstation is used to initialise and configure the test-bed, and to monitor the network during operation, for visualization and debugging purposes. However, it is not part of the sensor network and does not influence or control the system behaviour during normal operation. This workstation is the “System Visualization” block shown in Figure 3 (upper right), which is a schematic diagram of the multi-agent system architecture of the test-bed system.

A picture of the current state of the physical test-bed, with some panels removed to reveal the internal structure and electronics, is shown in Figure 2. A 12V power supply is mounted on the base of the hexagon, and power is distributed via the top and bottom edges. Communication between the test-bed and PCs is via 1.5 Mbits/s serial links using USB.

Figure 2. Photograph of the test-bed (top); a single cell consisting of a network application sub-module (NAS) and a data acquisition sub-module (DAS) (center photograph); the bottom photograph shows an aluminium panel containing four cells


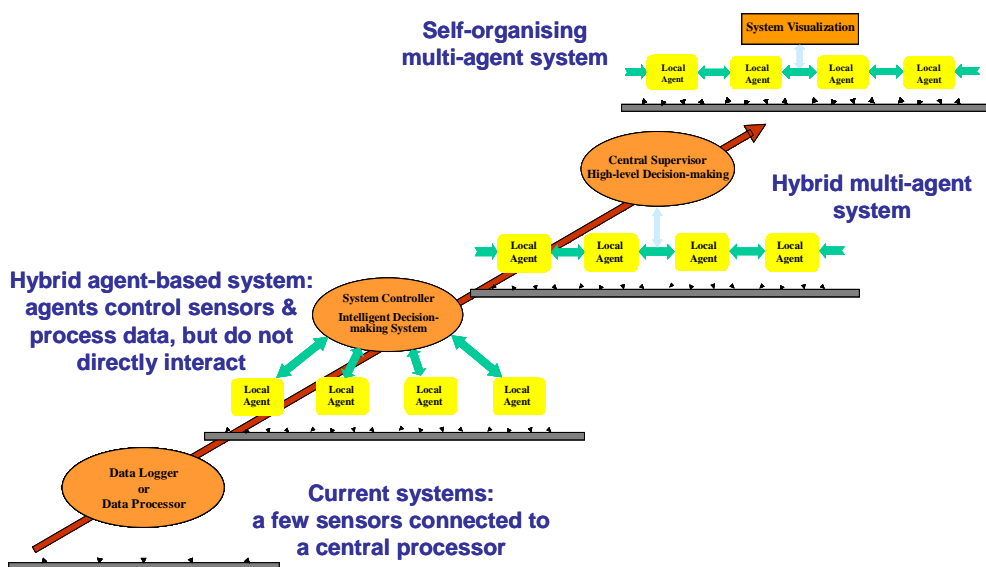


Figure 3. Schematic diagram of a progression of agent-based system architectures, leading to the complex multi-agent architecture of the test-bed system in the upper right of the diagram. Each cell in the test-bed is a local agent, capable of obtaining local information about damage from its sensors, and capable of communicating only with its direct neighbours in the mesh network. The workstation referred to in the text is the System Visualization block.



Cell Properties and Single-Cell Functionality

A modular approach to the sensing and electronics was adopted to enhance flexibility, re-configurability and ease of manufacture (Hedley et al., 2004, Batten et al., 2004). Each module, or cell, contains sensing elements, signal processing (analogue and digital), and communications, using the following logical layering of these functions:

1. **Sensors** — piezo-electric polymer sensors attached to the aluminium skin.
2. **Analogue signal processing** — including amplification, filtering, and other processing of the sensor signals required prior to digitization of the signals.
3. **Sampling and data pre-processing** — including digitization, calibration correction, data reduction, and other processing that can be performed using only the local signals.
4. **Data analysis and localised agent algorithms** — at this level, data is processed using information from local sensors and neighbouring modules.
5. **Inter-module communication** — comprising the software stack and the physical links to provide communication between modules.

These layers are divided into two groups, the data acquisition layer (DAL), which consists of layers 1 to 3, and the network application layer (NAL), which consists of layers 4 and 5. These are implemented as separate physical sub-modules, called the Data Acquisition Sub-module (DAS) and network application sub-module (NAS), respectively. This separation allows replacement of one type of sensor, and its associated electronics, with another sensor sub-module, without changing the main processing and communications hardware, hence allowing a range of sensor types to be tested.

The DAS provides gain and filtering for the four attached piezo-electric sensor signals (which have components up to 1.55 MHz after analogue filtering). These signals are digitized at 3 Msamples/sec using a 12-bit analogue-to-digital converter, and initial processing to estimate the time of arrival of a signal on each sensor is performed using digital signal processing. This information is passed to the NAS using a high-speed synchronous serial communications link, and power is received from the NAS over the same connector.

The NAS contains both a 400 MIPS fixed-point digital signal processor (DSP) and 400k gate field programmable gate array (FPGA), along with 2 Mbytes of non-volatile memory and 8 Mbytes of volatile memory. These resources are used by the software agents running in each cell, which provide the network intelligence. Each NAS contains a unique 64-bit identifier. An NAS contains four ports, used for communication with its four nearest neighbours and for power distribution. This provides a highly robust mesh network structure that will maintain connectivity even if a significant number of cells or communications links are damaged.

Impact Detection

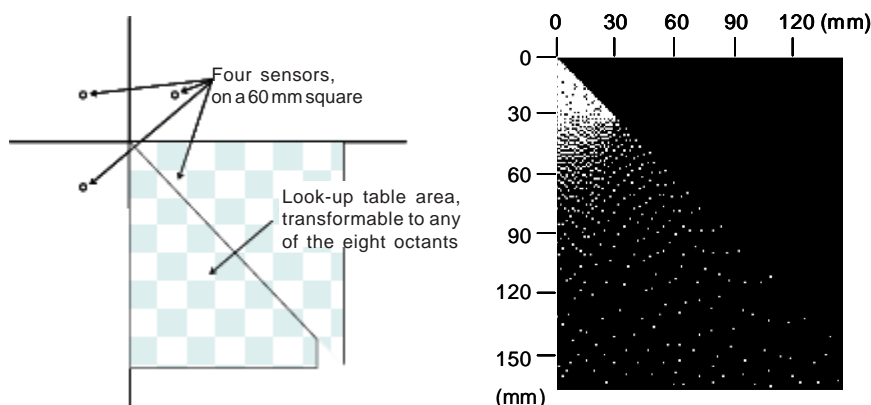
Piezo-electric sensors, consisting of a 2.5 mm-diameter, 110 μ m-thick film of PVDF (polyvinylidene fluoride) coated on both sides with a conductive gold layer, are bonded to aluminium sheets which form the “skin” of the concept demonstrator, providing a suitable method for detecting the plate waves in the aluminium sheets.

The DAS continuously samples four analogue channels, storing the data in a circular buffer containing 200 samples, or 64 μ s of data, from each channel, and checking if a sample has deviated by more than 90mV from the channel’s average value. Once a channel has exceeded this selected threshold, a further 184 samples are taken on each channel, an impact is flagged, and the buffers are processed.

The signals from the four sensors are detected by the DAL electronics, narrow-band filtered, amplified, and digitized. Then the *earliest arrival time* is subtracted from the other three times to give three time delays. These delays are used to estimate the location of the impact relative to the centre of the square formed by the four sensors used. In other words, the impact is triangulated based on measured arrival times of the lowest-order extensional wave and the known group velocity (about 5300 m/s) of these waves in the aluminium plate at a particular frequency (about 1.5 MHz).

If standard triangulation techniques were applied to these three time delays, the equations for three hyperbolae would need to be solved simultaneously. Further, as the digitization rate is limited to around 3 MHz in the present hardware, the time of arrival of the extensional waves may not be determined accurately enough for solutions to these equations to exist. Searching for near solutions can be done, but this would take a significant amount of processor time and memory in the present configuration.

Figure 4. Left — Position of sensors and the area of the panel covered by the look-up table (the table is transformed to the other eight octants to cover the full panel); Right — Graphical representation of look-up table points and estimates of the positional errors (the top left-hand corner represents the centre of the group of four sensors, and the axes' labels are millimetres from the centre of the square formed by the sensors. The white dots are the points in the look-up table: Their spacing is an indication of the uncertainty in the estimate of the location of an impact).



Instead, the present version of the demonstrator employed a faster way: a *look-up table*. As each cell has four sensors arranged at the corners of a 60 mm square, the look-up table will be the same for each of the eight octants whose origin is at the centre of the square. For the present size of aluminium panel, the largest area that needs to be covered by the look-up table is a truncated, isosceles (at 45 degrees) right triangle (Figure 4), 165 mm long and 143 mm in the truncated direction. This geometry will cover the “worst case” of an impact in a corner of the panel diagonally opposite that of the cell containing the four sensors being used. This area was then divided into 1-mm squares. For each square, the time delays may be calculated and used to form a six-figure “index” that is associated with each position. This index number is formed from the three delay times from each pair of sensors, the first two digits being the shortest delay time (in number of time step intervals of 320 ns), the second two-digit number is the next longest delay time, and the final two-digit number the longest delay time. While the total number of points (on a 1-mm grid) inside this truncated triangular area is about 13,600, the total number of unique points (points that have different index numbers) is only 1,071 due to the finite time-step interval. The mean position of all the 1-mm cells that have the same index number is stored in the look-up table with that index number. The look-up table is stored in the flash memory on the DSP.

When an impact is registered, the delays are first ordered to quickly determine which octant of the Cartesian plane contained the impact site. Depending on the order of the impacts (the time of arrival at each of the sensors) one of eight co-ordinate transformation

matrices is used to translate the look-up table result to the correct orientation. The index number of the impact is then found in the look-up table, and the average position associated with that index number multiplied by the appropriate transformation matrix is given as the location of the impact. In this way, any of the four sets of four sensors on a panel can detect the location of an impact.

DAMAGE SCENARIOS

Introduction

In the *Introduction* section we outlined the processes associated with a damage situation and categorised them in terms of a matrix whose two variables are spatial extent and degree of deliberation. This is a convenient means of describing a wide range of damage situations by breaking them down into components which range in extent from local to global and in the associated processing needed, from purely reactive (minimal processing) to strongly deliberative (requiring significant high-level cognition). This is similar to the response of biological organisms to damage, as was discussed.

Three different damage situations have been selected for discussion, two of them in considerable detail. These are (a) *critical damage*, defined as damage severe enough to threaten the integrity of the vehicle, generally caused by a sudden event (such as a major impact or explosion), (b) *sub-critical damage* which, although severe enough to require immediate action, does not invoke an emergency response, and (c) *non-critical* or *minor damage*, which does not necessarily require an immediate response, but which must be monitored continually in terms of its cumulative effects. Each of these will be described in terms of its components and analysed in terms of the two variables in the matrix, spatial extent and degree of deliberation. Since (a) is mostly reactive while (b) and (c) require varying degrees of deliberation, the latter two will be treated in sufficient detail to illustrate how the necessary computations may be accomplished in self-organising fashion in a multi-agent environment.

The following three sub-sections outline self-organizing responses to sub-critical and non-critical damage. This section, therefore, sets the context for these subsequent discussions.

Critical Damage

Critical damage means damage severe enough to threaten the integrity of the vehicle, requiring immediate action to ensure its survival. The initial response to such a situation must almost certainly be reactive, since the necessary actions will need to be implemented as rapidly as possible. It will almost certainly be global in extent, since the whole structure will need to know about events of this import, even if the initial reaction occurs in the neighbourhood of the damage site. In terms of the response matrix, a critical damage event may be represented as shown in Table 2.

In such situations time is of the essence, and both detection and communication must be done as quickly as possible. The requirements on the communications network are to send an alarm as rapidly as possible to one or more (probably) unknown locations

Table 2. Critical damage event response matrix

	Component	Spatial extent	Deliberation
	Detection	Local/Neighbourhood	Reactive
	Report	Neighbourhood or Global	Reactive
Actions:	Isolate damaged region	Global	Reactive

using only the adjacent-cell communication links. In addition, there may exist barriers to communication due to the network configuration itself or to significant prior and continuing damage. Thus the communications environment is largely unknown and changing, providing a major challenge. Some work has been done on these problems, but much more is needed (Li, Guo, & Poulton, 2004). Detection of critical damage by the cell network is also difficult because of the necessary time constraints, and it will almost certainly be good policy to err on the side of caution. Some kind of local activity measure will probably be best, but again very little research exists as yet.

The emergency response to critical damage will almost always be followed by more deliberative actions once the immediate safety of the vehicle has been assured. This is outlined in the response matrix, and follows a path very similar to that for sub-critical damage, which is described in the next sub-section.

Sub-Critical Damage

Sub-critical damage is taken to mean local damage to one or a number of cells which, although serious enough to require immediate action, does not threaten the immediate survival of the vehicle. In terms of the response matrix, such events may be broken down as described in Table 3.

Damage detection occurs at the local (cellular) level as described in the *Ageless Aerospace Vehicle* section, and is followed by a local response whose purpose is to define the extent of the damage and allow the assessment of its severity. For the AAV, this involves the self-organised formation of impact boundaries (Foreman, Prokopenko & Wang, 2003; Lovatt, Poulton, Price, Prokopenko, Valencia, & Wang, 2003; Prokopenko et al., 2005a; Wang & Prokopenko, 2004), which are described in some detail in the *Impact Boundaries* section below. When an assessment has been made, it must be communicated to some point on the vehicle from which appropriate action may be generated. This is not reactive, but may be at either neighbourhood or global level. The same issues apply regarding communications, which have been discussed in the *Critical Damage* sub-section.

The appropriate action will depend on circumstance, and three examples are given above. These range from local repair, which may indeed be reactive, to invoking a secondary inspection mechanism to obtain additional information about the nature and severity of the damage. This may be carried out within the neighbourhood, or it may require global interactions. An intriguing possibility for future aerospace vehicles is self-replication, where a replacement for the damaged section is manufactured by a self-organising process. This is discussed more fully in the *Shape Replication* section that follows.

Table 3. Sub-critical damage event response matrix

	Component	Spatial Extent	Deliberation
	Detection	Local	Reactive
	Local response	Neighbourhood	Moderately deliberative
	Assess	Neighbourhood	Moderately deliberative
	Report	Neighbourhood or Global	Strongly deliberative
Actions:	Secondary Inspection	Neighbourhood or Global	Strongly Deliberative
	Local repair	Local	Reactive or Moderately deliberative
	Self-replication	Global	Moderately deliberative

Non-Critical (Minor) Damage

The last damage scenario to be analysed in some depth is non-critical or minor damage. As outlined in the *Response Matrix* sub-section, such damage is typified by minor impacts, fatigue, or corrosion, processes which do not interfere immediately with the functioning of a cell, but which may lead to structural failure if accumulated over time. It is necessary to monitor such damage, not only to assess its long-term impact, but also because of its possible interaction with more serious types of damage. As would be expected considering the longer time-scale for this damage mechanism, assessment and action are quite deliberative, and of broader (neighbourhood or global) extent than previously discussed examples, although detection will most likely still be local. Referring to the response matrix, a possible breakdown for non-critical damage is as described in Table 4.

Although the detection of non-critical damage is local and reactive, all other steps are either broader in extent or degree of deliberation. This is to be expected since any assessment and resulting action must involve a number of non-critical damage sites. There is thus a need for the self-organisation of information regarding non-critical damage so that the relevant assessment can be made and acted upon. The most important of such information comprises the locations of damage sites and the severity of their damage, and there are several ways in which this information may be made to self-organise. One promising method is by the formation of an impact network, which is essentially a way of advantageously linking non-critical damage sites (Prokopenko et al., 2005a; Wang, Valencia, Prokopenko, Price, & Poulton, 2003). Not only does this make the necessary information available for processing, but it offers a mechanism by which secondary inspection (or repair) agents may rapidly assess the damage. The formation and use of impact networks is discussed in the section on *Impact Networks and Ant Colonies*.

IMPACT BOUNDARIES

Typically, the damage on the AAV skin caused by a high-velocity impact is most severe at the point of impact (an epicentre). It will be assumed that not only are the cells

Table 4. Non-critical damage event response matrix

	Component	Spatial extent	Deliberation
	Detection	Local	Reactive
	Local assessment	Local	Moderately deliberative
	Form impact network	Global	Moderately deliberative
	Report network status	Neighbourhood or Global	Moderately deliberative
	Assess network status	Global	Strongly deliberative
Actions:	Secondary inspection	Global	Strongly deliberative
	Repair	Global	Strongly deliberative

at the epicentre severely damaged, but that damage spreads to neighbouring cells, perhaps as a result of severe vibration, blast, or electromagnetic effects. One effect of this propagated damage is likely to be observed as damage to the communication capability of the neighbouring cells. For the sake of a specific example, we simulate the effect of this extended damage by assuming a communications error rate that is propagated out with an exponential decay to a certain radius (Lovatt et al., 2003). In this example, the damage can be characterised by a probability P_i of an error corrupting a message bit i , dependent on proximity of the affected communication port to the epicentre:


$$P_i = \frac{1}{2} \left(1 - \frac{d}{R}\right)^\alpha \quad (1)$$

where d is the distance between the involved communication port and the epicentre of the impact with the radius R , and α is the exponential decay of the communication loss (we have investigated a range of values, including linear decays, $\alpha = 1$, and high-order polynomial decays, $\alpha \leq 7$). Multiple impacts result in overlapping damaged regions with quite complex shapes.

In this section we describe multi-cellular *impact boundaries*, self-organising in the presence of cell failures and connectivity disruptions, and their use in damage evaluation and possibly self-repair.

It is desirable that an impact boundary, enclosing damaged areas, forms a continuously- connected closed circuit. On the one hand, this circuit may serve as a reliable communication pathway around the impact-surrounding region within which communications are compromised. Every cell on a continuously-connected closed circuit must always have two and only two neighbour cells, designated as the circuit members (circuit-neighbours of this cell). On the other hand, a continuously-connected closed impact boundary provides a template for repair of the impact-surrounding region, uniquely describing its shape (Figure 5). Both these functionalities of impact boundaries can be contrasted with non-continuous “guard walls” investigated by Durbeck and Macias (2002) that simply isolate faulty regions of the Cell Matrix, without connecting elements of a “guard wall” in a circuit. An impact boundary enables a shape-replication of a multi-

Figure 5. White cells are destroyed, red (dark-grey) cells form “scaffolding”, and blue (black) cells form the “frame” (boundary links are shown as white double-lines).



cellular impact-surrounding region, which can serve as an example of an end-to-end solution to an important SHM sub-task of damage evaluation and subsequent self-repair. The damage evaluation part (emergent evolvable impact boundaries) is implemented in the AAV-CD, while the shape-replication part is only simulated at this stage.

In order to serve either as a communication pathway or a repair template, an impact boundary should be *stable* despite communication failures caused by proximity to the epicentre, and such stability is our primary aim. In pursuing this aim, we deal with the following spatial self-organising layers:

- **Scaffolding** region, containing the cells that suffered significant communication damage;
- **Frame boundary**, an inner layer of normal cells that are able to communicate reliably among themselves; and
- **Closed impact boundary**, connecting the cells on the frame boundary into a continuous closed circuit by identifying their circuit-neighbours.

The “frame” separates the scaffolding region from the cells that are able to communicate to their normal functional capacity. In order to support a closed continuously-connected circuit, a *regular* frame should not be too thin (a scaffolding cell must not be adjacent to a normal cell), and should not be too thick (there must be no frame cells in the direction orthogonal to a local frame fragment). These internal layers (scaffolding, frame, and closed boundary) completely define an impact-surrounding region as a layered spatial hierarchy. In general, the impact-surrounding region can be seen as an example of annular spatial sorting: “forming a cluster of one class of objects and

surrounding it with annular bands of the other classes, each band containing objects of only one type” (Holland & Melhuish, 1999). It could be argued that, as an emergent structure, the impact-surrounding region has unique higher-order properties, such as having an *inside* and an *outside* (Prokopenko et al., 2005a).

Evolvable Localised Algorithm

The algorithm producing continuously-connected closed impact boundaries and the metrics quantitatively measuring their spatiotemporal stability are described in Foreman, Prokopenko, and Wang (2003) and Prokopenko et al. (2005a), while a genetic algorithm evolving agent properties which form impact boundaries satisfying these spatiotemporal metrics is presented in Wang and Prokopenko (2004). Here we briefly sketch the main elements of the evolved algorithm.

Every cell sends a *Ping* message to each of its neighbours regularly, and an *Acknowledgment* reply when it receives a *Ping* message. For each communication port i , a binary circular array A_i is used to store the communication histories for acknowledgments. The size of the array is called the communication history length μ . For each communication port, a Boolean success variable P_i is set to true if the percentage of *Acknowledgments* received in the A_i is greater than or equal to a certain threshold P . This variable is hysteretic: it changes only when a sufficient communication history is accumulated. This lagging of an effect behind its cause provides a temporary resistance to change and ensures a degree of stability in the treatment of communication connections between any two cells. A neighbour i is considered to be *communicating* when P_i is true. The algorithm uses the following main rules:

- Each cell switches to the *Scaffolding* state and stops transmitting messages if the number of communicating neighbours v is less than a certain threshold \mathcal{N} . For example, if $\mathcal{N} = 1$, then a cell switches to *Scaffolding* state if there are no communicating neighbours ($v < 1$).
- Each cell switches to *Frame* boundary state S_f if there is at least one communicating neighbour and at least one miscommunicating neighbour.
- Each cell switches to *Closed* boundary state S_c if the cell state is S_f , and there are at least two *communicating* neighbours.

The *Closed* boundary cells send and propagate (within a *time to live* period τ) specific *Connect* messages, leading to self-organisation of a continuous impact boundary. The cells that stopped transmitting messages may need to resume communications under certain conditions, for example, when a repair action is initiated, and their neighbours are again ready to receive communications (that is, when the cause of asymmetry is eliminated). The conditions for resumption of communications have to be precise so that they are not reacted upon prematurely, interfering with boundary formation. A variant of these *recovery* conditions is given:

- Each cell switches to the *Recovery* state S_r if the sequence P_0, \dots, P_3 describing the states of all four ports does not change for a specified number of consecutive cycles π_1 .
- A cell stays in the *Recovery* state S_r and may send communication messages during the next π_2 cycles.

In general, the described policy achieves the desired stability and continuity of self-organising impact boundaries. In addition, we have observed emergent spatiotemporal structures — recovery membranes — that separate the boundaries from recovering cells. A recovery membrane always forms on the inside of the closed boundary, and on the outside of the recovering area. Interestingly, unlike scaffolding and frame boundary, the membrane is not a designated state into which a cell can switch. Membrane cells shut down their communications like other scaffolding cells, but do not resume communications because *recovery* conditions are not applicable, as the miscommunicating neighbours are not stable. Without a membrane, the cells on the frame boundary would be confused by intermittent messages from scaffolding cells attempting recovery. Figure 6 illustrates a checkered-pattern recovery membrane shown with dark-grey colour, while the recovering cells are shown in yellow (darker shade of white).

The threshold \mathcal{N}_c limiting the number of communicating neighbours in switching to the *Scaffolding* state, significantly affects smoothness of a resulting boundary. In particular, if $\mathcal{N}_c = 1$ — in other words, a cell switches to the *Scaffolding* state if there are no *communicating* neighbours ($v < 1$) — then some boundary links may not be “smooth”: There are more than two ports connected by the link (Figure 7). If $\mathcal{N}_c = 2$ — that is, a cell switches to the *Scaffolding* state if there is at most one *communicating* neighbour ($v < 2$) — then all boundary links are “smooth” (for instance, the case depicted in Figure 6). If $\mathcal{N}_c = 3$, then any impact boundary is a rectangle. Finally, if $\mathcal{N}_c = 4$, then the impact-surrounding region fills the whole of the AAV array. This simple taxonomy of boundary types will be useful when we classify shape-replication algorithms as well.

Figure 6. Five white cells at the epicentre are destroyed, scaffolding cells that attempt recovery are shown in yellow (darker shade of white); a recovery membrane, shown in red (dark-grey), “absorbs” and separates them from the frame, shown in blue (black)





Figure 7. Four white cells are destroyed (a recovery membrane is not shown); the middle boundary link on the right-hand side is not smooth.



Spatiotemporal Stability, Phase Transitions, and Evolving Boundaries

The evolution of impact boundaries is based on spatiotemporal metrics incorporated within a fitness (objective) function. The analysis presented by Foreman, Prokopenko, and Wang (2003) and Prokopenko et al. (2005a) used two metrics to characterise stability of emergent impact boundaries: spatial and temporal.

The spatial metric is based on the variance in the size of the connected boundary-fragment (CBF). A CBF is simply a set F of cells in the *Closed* state S_c such that every cell in F is connected with at least one other cell in F , and there exists no cell outside F which is connected to at least one cell in F (an analogue of a maximally-connected sub-graph or a graph component). We calculate the maximum size $H_{sp}(t)$ of CBFs in self-organising impact boundaries at each cycle. Its variance σ_{sp}^2 over time is then used as a spatial metric within the objective function. This metric is inspired by random graph theory and is intended to capture *spatial connectivity* in impact boundaries. A continuous boundary may, however, change its shape over time, without breaking into fragments, while keeping the size of CBF almost constant. Therefore, a temporal metric may be required as well.

In order to analyse *temporal persistence*, we consider state changes in each cell at every time step. Given six symmetric boundary links possible in each square cell (“left-right”, “top-bottom”, “left-top”, and so on), there are 2^6 possible boundary states (including “no-boundary”), and $m = 2^{12}$ transitions. The entropy $H_{temp}(t)$ of a particular frequency distribution $S_i(t)$, where t is a time step, and i is a cell transition index: $1 \leq i \leq m$, can be calculated as follows (Equation 2):

$$H_{temp}(t) = -\sum_{i=1}^m \frac{S_i(t)}{n} \log \frac{S_i(t)}{n} \quad (2)$$

where n is the total number of cells, and $S_i(t)$ is the number of times the transition i was used at time t across all cells. The variance σ_{temp}^2 of the entropy $H_{temp}(t)$ over time is used as a temporal metric within the objective function.

Our task is complicated by the fact that emergent structures are characterised by a *phase transition* detectable by either σ_{sp}^2 or σ_{temp}^2 , rather than a particular value range. Therefore, simply rewarding low values for these entropy-based metrics would be insufficient. In particular, it has been observed (Foreman, Prokopenko, & Wang, 2003; Wang & Prokopenko, 2004) that both metrics are low-to-medium for algorithms with zero-length communication μ (tropicistic algorithms and chaotic regimes — Figure 8), increase dramatically for μ in the range $1 \leq \mu \leq \mu_0$, where μ_0 is a critical value at and below which complex unstable behaviours occur (Figure 9), and undergo a phase transition to very low values when $\mu > \mu_0$ (ordered regimes).

The critical value μ_0 is, of course, dependent on all other parameters used by the algorithm. Nevertheless, the chaotic regimes, which are more stable simply due to a small number of connections, can often be identified by a low average $\overline{H_{sp}}$ of the maximum sizes $H_{sp}(t)$ of CBFs in impact boundaries, ruling out at least zero-length histories. In particular,

Figure 8. A chaotic boundary with $\overline{H_{sp}} \leq 16$ and zero-length communication μ ; a membrane does not form at all (both σ_{sp}^2 and σ_{temp}^2 are low-to-medium)





Figure 9. An unstable boundary with μ close to its critical value; the membrane is fragmentary (both σ_{sp}^2 and σ_{temp}^2 are close to their peaks; i.e. a phase transition)



impact boundaries with the average $\overline{H_{sp}} \leq 16$ can be safely ruled out; the resulting chaotic patterns, illustrated in Figure 8, are of no interest.

On the other hand, a preference among ordered regimes towards shorter histories is another useful identifier of a phase transition and the critical value μ_0 . Besides, a shorter communication history μ enables a quicker response, as do lower values of τ and π_1 .

Thus, our experiments used minimisation of the following objective function:

$$f(\beta) = \begin{cases} M & \text{if } \overline{H_{sp}} \leq 16; \\ \frac{1}{2}(4.0\sigma_{sp}^2 + 10^5\sigma_{temp}^2) + \mu + \tau + \pi + \beta\overline{H_{sp}} & \text{if } \overline{H_{sp}} > 16 \end{cases} \quad (3)$$

where M is the maximal integer value provided by the compiler. The coefficient β reflects the relative importance of the length of impact boundaries in the objective function; sometimes it may be as important to obtain the smallest possible impact perimeter as it is to maintain the shortest possible communication history. We alternated between $\beta_1 = 0.25$ and $\beta_2 = 2.0$.

Each experiment involves an impact at a pre-defined cell and lasts 500 cycles; the first 30 cycles are excluded from the series $H_{sp}(t)$ and $H_{temp}(t)$ in order not to penalise

longer history lengths μ . We repeat the experiment three times for every chromosome (a combination of parameters) and average the objective (fitness) values obtained over these runs. The details of the genetic operators, the employed replacement strategy, and a comparative analysis of metrics are described in Wang and Prokopenko (2004). Here we only summarise the results.

The experiment minimising the objective function $f(0.25)$ evolved solutions with long robust and continuous impact boundaries with $\overline{H}_{sp} = 40$ (Figure 10), around large impact-surrounding regions, while requiring fairly short hysteresis: $\rho = 2$ and $\pi = 8$. The stabilisation of an impact boundary around a large region occurs at the periphery of the communication damage, where the communication failure probability falls to zero due to the error correction code, and the process has a cascading nature, where the boundary expands to eventually cover the entire impact-surrounding region. In summary, the case $\beta = 0.25$ results in longer boundaries that are sometimes capable of morphing without breaking into fragments.

On the other hand, minimisation of $f_{sp}(2.0)$ resulted in more compact impact-surrounding regions ($\overline{H}_{sp} = 32$, Figure 11) and thinner membranes, at the expense of longer hysteresis: $\rho = 6$ and $\pi = 8$. These boundaries generally keep the shape of a regular octagon. This case ($b = 2.0$) results in shorter boundaries that cannot morph without breaking into fragments, so any instability leads to fragmentation. Both solutions favoured $\tau = 1$ as expected for square cells.

Figure 10. A large checkered-pattern membrane, with short hysteresis, within a morphing but closed and continuous boundary ($\beta = 0.25$)





Figure 11. A small membrane, with long hysteresis, within a regular octagonal boundary ($\beta = 2.0$)



These results are promising and demonstrate the possibility for a multi-objective design of localised algorithms. In particular, the desired response time as well as size (and potentially, shape) of impact boundaries become the *design parameters* and may be specified in advance, leaving the precise logic and parameterisation of the localised algorithms to selection pressures. We believe that the proposed methodology is well suited to the *design at the edge of chaos*, where the design objective (for example, a specific shape) may be unstable, while other parameters (such as the response time) may be optimal.

The impact boundaries form patterns that may be used in damage assessment and diagnostics, as well as templates for repair, and provide reliable communication pathways around impact-surrounding regions. Their multiple roles illustrate two kinds of emergence: *pattern formation* and *intrinsic emergence*, distinguished by Crutchfield (1994):

- **Pattern formation** refers to an external observer who is able to recognise how certain unexpected features (patterns) “emerge” or “self-organise” during a process (for example, convective rolls in fluid flow, and spiral waves and Turing patterns in oscillating chemical reactions). The patterns may not necessarily have specific meaning *within the system*, but obtain a special meaning *to the observer* when detected;

- **Intrinsic emergence** refers to the emergent features which are important *within the system* because they confer additional functionality to the system itself, like supporting global coordination and computation (for instance, the emergence of coordinated behaviour in a flock of birds allows efficient global information processing through local interaction, which benefits individual agents).

In the next section we shall illustrate how stable and continuous impact boundaries can be used as templates for multi-cellular shape-replication, while the following section will describe a self-organising communication mechanism among remote cells. This mechanism may be used, in particular, to communicate the information represented by the emergent boundary patterns to remote cells playing the role of observers and/or controllers, if necessary.


SHAPE REPLICATION: TOWARDS SELF-REPAIR

In general, individual failed cells can be replaced one by one — this is, after all, the point of having a scalable solution. Moreover, any impact-surrounding region enclosed within its impact boundary can also be repaired by replacing individual failed cells one by one. Sometimes, however, it may be required to replace an impact-surrounding region in one step, for example, to minimise the overhead of disconnecting individual cell-to-cell links. Replacing the whole region within a boundary would require a removal of only the links between the boundary and normal cells. In this sub-section, we provide an example where a self-organised impact boundary (produced by the evolved algorithm) may be used in self-repair, or more precisely, in shape replication.

Given the planar grid topology, each cell on the closed impact boundary may have six boundary links, connecting ports “left-right”, “left-top”, and so on. Enumerating four communication ports from zero to three (“bottom” to “right” clockwise) allows us to uniquely label each boundary link with a two-digit number λ — for example, “32” would encode a link between the “right” and “top” ports (Figure 12). Then, the whole impact boundary can be encoded in an ordered list of these labels. However, in order to replicate the bounded shape, filling it cell by cell, we need to introduce a system of coordinates relative to a cell containing the shape list. More precisely, the boundary genome is a list of triples (α, β, λ) , where (α, β) are relative coordinates of a cell with the boundary link λ .

The shape replication algorithms developed in the context of AAV (Prokopenko & Wang, 2004) are based on the principles of multi-cellular organisation, cellular differentiation, and cellular division — similar to the embryonics approach (Mange, Sanchez, Stauffer, Tempesti, Marchal, & Piguet, 1998; Sipper, Mange, & Stauffer, 1997). A desired shape is encoded when an emergent impact boundary inspects itself and stores the “genome” in a “mother” cell. The genome contains both data describing the boundary and a program of how to interpret these data. The mother cell is then seeded in a new place outside the affected AAV array. An execution of its program initiates *cell-replication* in the directions encoded in the genome (Figure 13). Each cell-replication step involves copying of the genome (both data and the program) followed by differentiation of the data: an appropriate shift of certain coordinates. Newly produced cells are capable of


Figure 12. Boundary links




cellular division, continuing the process until the encoded shape is constructed (Figure 14).

Two algorithms for the AAV shape replication are described by Prokopenko and Wang (2004). The first algorithm solves the problem for connected and disconnected shapes. The second algorithm, in addition, recovers from possible errors in the “genome”, approximating missing fragments. In particular, the genome is partially repaired (Figure 15) within each cell which detected a discontinuity. Although the repaired genome does not cover all the missing cells, it does not introduce any cells which were not in the original shape, exhibiting the soundness but not completeness property. In other words, the repaired boundary is contained within the original shape. Importantly, there is a redundancy in the shape replication process: other cells which did not suffer any damage would successfully replicate the parts not encoded in the partially repaired genomes.

Figure 13. Shape replication: Boundary cells encoded in the genome but not yet produced are shown with dashed lines.




Left: A black cell (seed) produces two white cells, indicated by arrows. Right: Two more cells are being produced: one of them is a scaffolding cell, pointed to by the horizontal arrow (the inside direction is recognised by the vertical strip being ‘crossed’ above and below the considered location).

Figure 14. Completed shape replication

The shape-replication algorithms handle both standard (“blueprint”) and non-standard emergent shapes, self-organising in response to damage. Moreover, it is possible to combine these types. For example, structural data can be encoded in the form of triples, and a given genome can be extended in run-time with the data produced by self-inspecting emergent boundaries. Similarly, the self-repair phase within a cell which detected an anomaly in the genome may draw some data from the structural “blueprints” rather than approximate segments between disconnected fragments.

Importantly, the first algorithm, not involving a recovery of the genome from possible errors, replicates shapes encoded in either smooth or non-smooth boundaries; it does not depend on the threshold \mathcal{N} limiting the number of communicating neighbours in switching to the scaffolding state. The second, genome-repairing, algorithm, however,

Figure 15. The cell shown inside a circle attempts self-repair

Left: the corrupted triples are shown with the ‘star’-like signs. Right: the repaired triples are marked with crosses.

cannot deal with non-smooth boundaries — that is, when $\mathcal{N} = 1$ and a cell switches to the *scaffolding* state if there are no communicating neighbours. Thus, adding the selection force rewarding genome-recoverability would lead to evolution of only smooth and stable impact boundaries ($\mathcal{N} = 2$). In other words, the taxonomy of boundary types, based on the threshold \mathcal{N} , is related to a classification of shape-replication algorithms; for example, it is conceivable that some replication sub-tasks may tolerate only rectangular shapes ($\mathcal{N} = 3$).

The shape replication process described in this section can be used in repairing the impact-surrounding regions in one step. As mentioned before, this is not the only feasible strategy, and failed cells can be replaced individually. In addition, there is a possibility to employ self-healing materials; however, this reaches beyond the scope of our investigation.

IMPACT NETWORKS AND ANT COLONIES

Decentralised inspection across the AAV network array may require an *impact network* among cells that registered impacts with energies within a certain band (for example, non-critical impacts). The self-organising impact networks create an adaptive topology allowing inspection agents (communication packets or, potentially, swarming robots) to quickly explore the area and evaluate the damage (for example, identify densities of impacts typical for a meteor shower, evaluate progression of corrosion, or to trace cracks propagation) — particularly where a number of individually non-critical damage sites may collectively lead to a more serious problem. Robotic agents may need an impact network which solves a travelling salesperson problem (TSP). On the other hand, a shortest or *minimum spanning tree* (MST) is often required in order to enable decentralised inspections when virtual (software) agents are employed, and may provide a useful input for the TSP. In this section we present an extension of an ant colony optimisation (ACO) algorithm, using an adaptive dead reckoning scheme (ADRS) and producing robust and reconfigurable minimum spanning trees connecting autonomous AAV cells. A novel heuristic is introduced to solve the blocking problem: reconfiguration of an existing path which is no longer available or optimal. Dynamic formation of a robust reconfigurable network connecting remote AAV cells that belong to a specific class was analysed in our previous work (Abbott et al., 2003; Wang et al., 2003). The ACO algorithm developed in these studies successfully approximates minimum spanning trees, but occasional alternative paths around critically damaged areas may still emerge, competing with the shortest paths and slowing the algorithm's convergence.

Let us define an AAV impact network. A two-dimensional AAV array can be represented by a planar grid graph $G(V, E)$: the product of path graphs on m and n vertices, which are points on the planar integer lattice, connected by the edges $E(G)$ at unit distances (Figure 16). The cells which represent specific points of interest (for example, the cells which detected non-critical impacts, or the cells playing a role of local “hierarchs”, “observers”, or “controllers”) form a sub-set P of $V(G)$. We need to identify those edges Z in $E(G)$ which connect the vertices in P minimally, so that the total distance (a sum of unit distances assigned to edges Z) is shortest. This problem is essentially the standard minimum spanning tree problem, except that a spanning tree is defined for a


graph, and not for a set of vertices. Our problem is sometimes referred to in the literature as the *rectilinear minimum (terminal-) spanning tree* (RMST) problem — a fundamental problem in VLSI design — and in which the vertices in P are referred to as *terminals* (Kahng & Robins, 1995). The important difference between MST and RMST is that rather than choosing MST edges out of the graph edges $E(G)$ directly connecting pairs of vertices, we need to find multi-edge rectilinear paths between vertices in P , minimising the total distance. This can be done via an auxiliary complete graph A , whose vertex set is P and in which the edge pq for $p, q \in P$ with $p \neq q$ has length equal to the Manhattan distance between nodes p and q . After a standard MST A_m is identified in the graph A , we merely need to convert all edges in A_m to rectilinear paths on the grid graph G (Figure 16).

However, the impact network problem, however, is complicated by possible “obstacles” created by discontinuities in the AAV grid graph G . Initially, the grid graph G is *solid*; it does not have any “holes”, so its complement in the infinite orthogonal planar grid is connected. New critical impacts may create such holes in the grid. Figure 16 illustrates the RMST problem with two scenarios. The first case is shown in the top part, and involves three edges and a simple MST A_m with the total distance of eight. The second case is shown in the bottom part: some cells are destroyed (the corresponding vertices are removed), and the auxiliary complete graph should be updated because one shortest path has changed (from five to seven). This requires a recomputation of its MST (the new MST distance is nine), with another edge being selected and converted to a rectilinear path. This illustrates that a new obstacle may not just require that a new shortest path is found between the two involved cells (the problem investigated by Wu, Widmayer, Schlag, & Wong, (1987)), but rather than the whole MST is re-evaluated. Moreover, there are cases when a cell/terminal is no longer needed to be included in the RMST, or a new cell/terminal needs to be added. Incremental updates of an old rectilinear spanning tree may provide a practical solution, but a quick divergence from an RMST is a significant problem.

Thus, from a graph-theoretic standpoint, the representation of the impact network problem changes over time due to insertion of new nodes (for example, non-critical impacts) or deletion of old nodes no longer fitting the impact range, while the problem’s properties change due to varying connection costs (for example, critical impacts destroying existing paths). In short, the problem changes concurrently with the problem-solving process (Prokopenko et al., 2005a), and we need a dynamic and decentralised computation of a rectilinear minimum terminal-spanning tree in the presence of obstacles. If the information (such as the auxiliary graph A) was available in one central point, then the RMST problem would essentially become an MST problem, with a subsequent conversion to rectilinear paths. In this case the required computation itself would not be NP-hard, although the fully dynamic case, in which both insertions and deletions must be handled online, without knowing the sequence of events in advance, would still be quite intensive. Eppstein (1996) estimated a running time of a fully dynamic graph minimum spanning tree algorithm as $O(n^{1/2} \log^2 n + n^\epsilon)$, where ϵ is a (very small) constant, *per update*. In our case, the auxiliary graph A is not even known at any single node/cell, so the desired algorithm should be both *decentralised* and *fully dynamic*.

These factors suggested that the problem of forming minimum spanning trees on the AAV skin can be efficiently tackled by ant colony optimisation (ACO) algorithms,

Figure 16. Three impact nodes are shown in black



The top-left figure shows a complete auxiliary graph A (dashed lines) with three edges. The conversion of its MST to rectilinear paths on the AAV grid graph is shown in the top-right figure (bold edges). Two lower figures show the graph with some vertices removed. The bottom-left figure shows an updated auxiliary graph A, and the bottom-right figure shows conversion of the new MST to rectilinear AAV paths.

proposed and enhanced over recent years by Dorigo and his colleagues (Coloni, Dorigo, & Maniezzo, 1992; Dorigo & Di Caro, 1999; Dorigo, Maniezzo, & Coloni, 1996), rather than distributed dynamic programming (Bellman-Ford) algorithms. Essentially, the ACO algorithms use the ability of agents to indirectly interact through changes in their environment (*stigmergy*) by depositing pheromones and forming a pheromone trail. They also employ a form of *autocatalytic* behaviour — *allelomimesis*: the probability with which an ant chooses a trail increases with the number of ants that chose the same path in the past. The process is thus characterised by a positive feedback loop (Dorigo, Maniezzo, & Coloni, 1996). An overview of the ACO meta-heuristic and its applicability can be found in Dorigo and Di Caro (1999).

In the AAV-CD the ants are implemented as communication packets, so the policies are implemented via appropriate message passing, where the cells are responsible for unpacking the packets, interpreting them, and sending updated packets further if necessary. Thus, ants cannot move into the cells with damaged (or shutdown) communication links, so critically-impacted cells form obstacles, and the ants are supposed to find the shortest paths around them using positively-reinforced pheromone trails. For our problem, it is impractical to use two types of pheromone (such as “nest” and “food”) because each impact cell (node) serves both as a “nest” and a “food” source. Therefore, having two types of pheromone per node would have created multiple pheromone fields, combinatorially complicating the network. In addition, dissipation of pheromone over

large distances is not practical either, as it would lead to “flooding” of the network with messages. Hence, the algorithms developed for the AAV network use only one type of non-dissipative evaporating pheromone.

The ACO-ADRS Algorithm

The algorithm presented in Abbott et al. (2003) and Wang et al. (2003) was based on a hybrid method of establishing impact networks, using a single impact gradient field (IGF) and a dead reckoning scheme (DRS), complementing the autocatalytic process of ant-like agents. Following Prokopenko et al. (2005a) and Prokopenko, Wang, and Price (2005), we summarise here a major variant of this algorithm, without an IGF, and relying only on DRS. The behaviour of exploring ants includes the following:

- (E1) each impact node generates a number of exploring ants every T cycles; each ant has a “time to live” counter t_k , decremented every cycle;
- (E2) an exploring ant performs a random walk until either (a) another impact node is found, or (b) the ant has returned to the home impact node, or (c) the ant can move to a cell with a non-zero trail intensity;
- (E3) if an exploring ant can move to a cell with a non-zero trail intensity, the destination cell is selected according to transitional probabilities;
- (E4) at each step from cell i to cell j , an exploring ant updates the x - and y -shift coordinates from the home node (initially set to 0).

The DRS requires that each ant remembers the x - and y -shift coordinates from the home node. These coordinates are relative, they simply reflect how many cells separate the ant from the home node in terms of x and y at the moment, and should not be confused with a “tabu” list of an ACO agent containing all visited nodes in terms of some absolute coordinate or identification system. The DRS enables the agents to head home when another impact node is located:

- (R1) when another impact node is found, the exploring ant switches to a return state, remembers the ratio $g = y/x$ corresponding to the found node’s coordinates relative to the home node, and starts moving back to the home node by moving to cells where the y - and/or x -shift coordinates(s) would be smaller and their ratio would be as close as possible to g ; if both x - and y -shift are zero (the home node), the returning ant stops;
- (R2) if the cell suggested by the DRS (minimisation of x - and/or y -shift, while maintaining g) cannot be reached because of a communication failure (an obstacle), the ant selects an obstacle-avoiding move according to the transitional probabilities; upon this selection the ant keeps to the chosen path until the obstacle is avoided, as recognised by comparison of current y/x ratio with g ;
- (R3) each cycle, a returning ant deposits pheromone in the quantity inversely proportional to the traversed return distance q (q is incremented by 1 each cycle); the deposited pheromone is limited by a pre-defined maximum ϕ_{\max} .

The pheromone is deposited on the cells themselves rather than communication links; we deal with pheromone trail intensities ϕ_j at the cell j , used in calculating transitional probabilities and determining which neighbour cell should be chosen by an incoming ant packet to continue their travel. The intensity of trail $\phi_j(t)$ on the node j gives

information on how many ants have traversed the node in the past, and is updated each time an ant agent k passes through the node:

$$\varphi_j(t) = \min(\varphi_j(t) + \frac{\sigma_k}{q_k(t)}, \varphi_{\max}) \quad (4)$$

where σ_k is a constant quantity specified for each generated ant k , q_k is the distance traversed by the ant k , and φ_{\max} is a limit on pheromone trail intensity. Intuitively, the quantity σ_k represents a pheromone reserve of the ant k , consumed during the return trip. At the beginning of each cycle, the pheromone evaporates at the rate $\rho \in (0,1)$:

$$\varphi_j(t) = (1 - \rho) \varphi_j(t) = \psi \varphi_j(t) \quad (5)$$

where ψ is the retention rate. A study of the impact network stability is provided by (Prokopenko et al., 2005a). An improvement to the DRS algorithm included *adaptive* pheromone reserve quantity σ_k and “time to live” counter τ_k , and a “pause” heuristic (Prokopenko, Wang, Scott, Gerasimov, Hoschke, & Price, 2005b). The pheromone reserve is adaptively allocated by the generating node, based on the ants returned to the node in the past:

$$\sigma_k = \max(\gamma_1 \hat{q}, \sigma_{\min}) \quad (6)$$

where \hat{q} is the minimal distance traversed by the returned ants, γ_1 is a scaling factor, and σ_{\min} is a lower limit for the pheromone reserve allocated for an ant. Analogously,

$$\tau_k = \min(\gamma_2 \hat{q}, \tau_{\max}) \quad (7)$$

where τ_{\max} is an upper limit for the counter, and γ_2 is a scaling factor. Equations (6) and (7) define the adaptive dead reckoning scheme (ADRS), which contributes to a faster reconfiguration of trails and minimum spanning trees.

The “pause” heuristic contributes to a better convergence of the DRS and ADRS algorithms. Let us consider decisions of a returning ant in the situation when an obstacle blocks a DRS path towards the home node. If the ants used both “nest” and “food” pheromones, then an ant returning to the “nest” and tracing the “nest” pheromone (while depositing the “food” pheromone) would benefit from the stigmergy as both shorter and longer paths around the obstacle were chosen in the past by some ants going in the opposite direction. In other words, when an ant is at the “decision” node, the transitional probabilities would reflect the difference between the alternative trail-to-nest intensities. Similarly, an ant tracing the “food” pheromone would use at the decision node the difference between the alternative trail-to-food intensities created by the returning ants that have traversed either shorter or longer paths around the obstacle in the past. This feature is very important in the beginning — when a new obstacle appears — and the transitional probabilities at the decision node are uniformly distributed. The autocatalytic process is then “kick-started” by the ants going in the opposite direction and using a different pheromone type. The ants going along a shorter return path deposit more pheromone than the ants that select a longer path around an obstacle — simply because

the deposited quantity is inversely proportional to the traversed distance. A higher quantity of pheromone attracts more ants. Eventually, the alternative shortest path is established between a pair of impact nodes.


The DRS algorithm uses only one type of pheromone (the “impact” pheromone). Therefore, the ants going in the opposite directions and using the same pheromone type obscure the difference between the alternative trail intensities at the decision node, confounding the choice. For example, a returning ant facing an obstacle ahead and excluding a backtrack possibility has a 50:50 chance of turning left or right, when the trails are not yet established. Choosing a direction at this decision node results in the ant depositing the pheromone either on the left or the right node. Clearly, this deposit is not an informed choice, being driven by a 50:50 chance, and may in fact obscure the pheromone trail. The update of the pheromone on both left and right nodes should, in fact, be done *only* by the ants going in the opposite direction, as these ants have traversed an alternative path. To reiterate, this dilemma is not present when the ants use two types of pheromones. A simple solution enhancing the DRS algorithm, using only one pheromone type, is provided by the “pause” heuristic:

(R4) an ant, facing an obstacle at cycle t and making a transition to the next node, does not deposit any pheromone at cycle $t+1$, resuming pheromone deposits only from cycle $t+2$.

The “pause” heuristic initially produces gaps in the trails, next to each decision point (Figure 17). However, these gaps are eventually filled by the ants going in the opposite direction, leading to the reinforcement of the shortest trail. Figures 17-19 illustrate this dynamic with snapshots of the 24 x 8 AAV-CD network array, visualised by the Debugger tool.

The enhanced ADRS algorithm produces rectilinear minimum spanning trees, resulting in reconfigurable impact networks, and performs well in dealing with two well-known problems: the blocking problem and the shortcut problem. *Blocking* occurs when a trail that was found by the ants is no longer available due to an obstacle and an

Figure 17. White cells detected non-critical impacts



An initial vertical trail is destroyed by a horizontal obstacle (seven cells are removed). The returning ants explore two alternative possibilities. The gaps in both trails form next to each decision node.

alternative trail is needed. The *shortcut* corresponds to a new shorter trail becoming available due to repaired cells (Schoonderwoerd, Holland, Brutton, & Rothkrantz, 1997).

Experimental Results

The analysis of algorithm convergence is based on the concept of a connected trail-fragment (CTF). A CTF is a set F of cells with $\phi \geq \vartheta$ (where ϑ is a given threshold), such that every cell in F is connected with at least one other cell in F , and there exists no cell outside F which is connected to at least one cell in F . We focus here on one important design parameter: pheromone retention rate ψ which determines how much pheromone is left in the cell at the end of each cycle ($\psi = 1.0$ means that there is no evaporation). We carried out 10 experiments with three impacts, for different pheromone retention rates between 0.1 and 0.99. During each experiment, we calculated the average size of CTFs in impact networks, $H(\psi)$, at each time-point, and its standard deviation, $s(\psi)$, over time. It was observed that low retention rates (for instance, $\psi = 0.86$) lead to chaotic trails; critical retention rates (such as, $\psi = 0.94$) lead to unstable trails (“the edge of chaos”); and high retention rates (for example, $\psi = 0.98$) support stable trails. The critical retention rates between $\psi = 0.90$ and $\psi = 0.94$ result in the most “complex” dynamics: a trail frequently

Figure 18. The gaps of the shorter trail are filled, while the longer trail slowly evaporates without gaps being robustly filled





Figure 19. A shorter trail is established



forms and breaks. A detailed picture emerges from Figures 20 and 21. Both plots indicate a clear peak in the standard deviation: $\sigma(\psi)$ peaks in the neighbourhood of $\psi = 0.9$, making the phase transition apparent, and clearly separating ordered and desirable robust scenarios from chaotic and under-performing cases.

Thus, tracing the average size H of CTFs, and its variance σ^2 , over time allowed us to identify emergence of rectilinear minimum spanning trees as a phase transition in network connectivity. Our experiments have shown that the ADRS algorithm enhanced with the “pause” heuristic outperforms the original algorithm in terms of these metrics. In particular, we first compared the performance of these two variants in a scenario without obstacles, focusing on the contribution of *adaptive* pheromone reserve quantity and time-to-live counter. The comparison between maximums of $\sigma(\psi)$ for the enhanced ADRS algorithm and the original variant, where the latter was evaluated over three experiments, shows that the ADRS enhancement results in approximately 9% less dispersed data at the edge of chaos (in other words, the standard deviation at its maximum is 2.57 against 2.79, given the same mean size of CTFs), and a more pronounced minimum of $\sigma(\psi)$ after the phase transition, at the retention rate $\psi = 0.96$.

Secondly, we compared the algorithms in a scenario with two impacts and an obstacle, focusing on the contribution of the “pause” heuristics to the solution of the blocking problem. We carried out 10 experiments for each value of the pheromone retention rate in the range between 0.81 and 0.99. During each experiment, a simple straight trail (length 9) was initially formed between two obstacles, and then broken at cycle 200. As before, we calculated the average size of CTFs in impact networks, $H(\psi)$, at each time-point, and its standard deviation, $\sigma(\psi)$, over time. The same three types of dynamics, chaotic, complex, and ordered, were observed. This scenario is more challenging because two “ordered” phases are observed (Figure 21). The first (and the one we are interested in) is the emergence of the stable shorter trail around the obstacle (length 15) as opposed to the longer trail (length 21), followed by the emergence of both stable trails around the obstacle (combined length 29). The first “ordered” phase is separated from the chaotic phase ($\psi < 0.94$) by the “edge of chaos” ($\psi = 0.94\text{--}0.96$), and is identified by the minimum of $\sigma(\psi)$, also at the retention rate $\psi = 0.96$. The second “ordered” phase occurs at very high retention rates $\psi \geq 0.99$, and is of no interest: there is enough pheromone to support many trails.

Thus, in terms of solving the blocking problem, the optimal pheromone retention rate ψ can be identified as the one which attains the minimum of the standard deviation $\sigma(\psi)$, following the edge of chaos pointed to by the first maximum of $\sigma(\psi)$, as we increase ψ . When the optimal rate ψ is identified, one can compare the performance of the algorithms at their optima.

The algorithm enhanced with the “pause” heuristics is as good as the main variant in terms of the time it takes for the shorter trail to become the primary choice (on average 147 cycles after the obstacle, for the new algorithm, against 152 cycles for the main variant), and significantly outperformed it in terms of data dispersion both at the edge of chaos and at the optimum:

- the average (over 10 runs) standard deviation at its first edge-of-chaos maximum is 4.92 against 6.08 (24% improvement), and
- the average standard deviation at its first ordered-phase minimum is 3.48 against 5.01 (44% improvement), given the same mean size of CTFs.

Importantly, shorter trails around the obstacle appear as quickly as before but are much more stable with the modified algorithm. In other words, the pause heuristic does not delay emergence of the shorter trail as the primary choice, but makes resultant trails significantly more stable. The main share of the improvements is due to the “pause” heuristics rather than ADRS (which improves dispersion in the order of 10%).

In this section, we considered the emergence of impact network pre-optimising decentralised inspections on an AAV skin, and introduced a new local heuristic improving performance of the modified ACO-DRS algorithm. In summary, the modified algorithm involves one type of non-dissipative evaporating pheromone, simple ant-routing tables containing normalised pheromone values only for immediate neighbours, one type of ant with limited private memory; the dead reckoning scheme, and the transitional probabilities model with obstacle threshold. The ADRS algorithm enhanced with the “pause” heuristic is deployed in the AAV-CD and robustly solves blocking and shortcut problems, producing rectilinear minimum spanning trees for impact-sensing networks.

Figure 20. Background (left z-axis): average size H of CBF's, for different retention rates: $\psi < 0.90$ (chaotic), $\psi = 0.90\text{--}0.94$ (unstable), and $\psi > 0.94$ (stable). Foreground (right z-axis, in red colour): standard deviation $\sigma(\psi)$ of the average size $H(\psi)$




Figure 21. The scenario with three impacts and no obstacles. Standard deviation s of the average size H , for the enhanced ADRS algorithm; phase transition is evident in the range $\psi = 0.90-0.94$





Figure 22. The scenario with two impacts and an obstacle. A chaotic phase (fragmentary trails) is separated by the edge of chaos (first maximum is at $\psi = 0.94$) from the first ordered phase (stable short trails), followed by another phase transition to combined trails ($\psi > 0.98$)



While we have not evolved parameters for the ACO-DRS algorithm, the observed phase transitions clearly identify the critical values that would be selected by a genetic algorithm (GA) in rewarding stable pheromone trails — similarly to the evolution of impact boundaries guided by the stability metrics. In any case, an optimisation technique should provide a very comprehensive exploration at the edge of chaos. For example, the first phase transition in the dynamics produced by the main variant of our algorithm is in the range $\psi = 0.975\text{--}0.977$, and can be easily missed by a GA with inadequate replacement strategies.

DISCUSSION: SELF-ORGANIZATION SELECTION PRESSURES

Self-organising solutions presented in the preceding three sections depend on selection pressures or forces which, through their contribution to the evolutionary fitness functions, constrain the emergent behaviour. One example of a generic selection pressure is the *spatiotemporal stability* of emergent patterns: arguably, any pattern has to be stable before exhibiting another useful task-oriented feature. The sub-critical damage scenario illustrated the use of spatiotemporal stability in evolving impact boundaries (*Impact Boundaries* and *Shape Replication* sections). The impact networks (*Impact Networks and Ant Colonies* section) employ stability as well: the observed phase transitions clearly identify the critical values that would satisfy a fitness function rewarding stable pheromone trails.

Another example of an independent selection force is network *connectivity*, which rewards specific multi-agent network topologies. This force, we believe, is related to both efficiency and robustness, which were identified by Venkatasubramanian, Katare, Patkar, and Mu, (2004) as critical measures underlying optimal network structures. In this context, the efficiency of a graph is defined as the inverse of its average vertex-vertex distance, and is related to the short-term survival. Effective accessibility is defined via a number of vertices reachable from any vertex of a graph component, added over all components. Intuitively, it identifies how quickly a vertex can be reached from other vertices. Structural robustness is then defined with respect to a vertex as the ratio of the effective accessibility of the graph, obtained by deleting this vertex from the original graph, to the maximum possible effective accessibility. Intuitively, this measure captures the importance of this vertex to the connectivity (accessibility) of the graph; in other words, how much the connectivity would be affected if this vertex is removed. Using these definitions, it is possible to define average-case structural robustness as the average computed over all the vertices, or worst-case structural robustness as the minimum computed over all the vertices. Venkatasubramanian et al. (2004) argue that after removal of a vertex, some or all of the sub-graphs could still be functional, and use normalised efficiency of the largest remaining component as an indicator of the functional robustness of the system after damage, relating it to the long-term survival. Both efficiency and robustness identify, in our view, aspects of *connectivity* needed for emergence of optimal multi-agent networks.

Another important selection force is an information-storage ability and *self-referentiality* of representation, providing an emergent pattern with a means for replica-

tion. A well-known example is crystal growth, involving template-based copying process and error correction, and preserving aspects of the crystal structure on macroscopic scales. Each crystal stores a self-referential template (for example, the cross-sectional shape), which may be used in reproduction by splitting. This self-referring arrangement is arguably very simple: the template for growth is the crystal's cross-section, which is directly used in the crystal growth. A much more involved example is the genotype-phenotype relationship, where the degree of self-referentiality is much higher, and the reproduction involves many intermediate steps, mapping genotype into phenotype.

The shape replication process described in the *Shape Replication* section can also be explained in self-referential terms (Prokopenko & Wang, 2004), employing two logical levels. It is well-known that self-replication of a system can be characterised by emergent behaviour and *tangled hierarchies* exhibiting Strange Loops:

an interaction between levels in which the top level reaches back down towards the bottom level and influences it, while at the same time being itself determined by the bottom level. (Hofstadter, 1989)

The shape replication process can be described in these terms as well. An impact boundary emerges at a level which is higher than the object level where individual cells are interacting. The genome of the enclosed multi-cellular shape is a model of an impact boundary, and embedding this higher-level model within every involved cell at the object level is self-referential.

The genome model is obtained by self-inspection of the impact boundary. The process of self-inspection is mirrored by the self-inspection of the genome, carried out internally by each cell at every replication step in order to detect discontinuities in the encoded boundary. Similarly, self-repair of the entire damaged impact-surrounding region is reflected in the internal self-repair of the model (genome). Following Hofstadter's language, the top-level pattern (a boundary) emerges itself out of interactions of cells, while also reaching down to the bottom level and influencing it. This example with self-referential shape replication did not involve explicit metrics for self-referential inspection and repair processes. Nevertheless, our conjecture is that the degree of self-referentiality can be measured and used in evolving multi-agent networks.

Responses to critical damage highlight the role of a low computation and communication complexity as another selection force. The main principle in considering emergency and/or "panic" responses is that the system needs to alter its priorities from long-term survival to emergency short-term survival, on many levels. In terms of the AAV, an emergency response may therefore require changing priorities of communication messages, an increase in the rate of polling the buffers of the communication ports, redirection of more power to specific modules, while temporarily disabling other modules, and so on. Subsequently, it may cause an activation of secondary passive and mobile sensors. This cascading scenario requires a fast and unconscious (un-reasoned) reaction, immediately upon a detection of a specific sensory input (trigger). This trigger should be detected locally, simply because detecting and matching near-simultaneous remote sensory inputs would have to be done "deeper" within the system, leaving less time for the emergency response. In other words, the trigger is "locally-situated" both in space and time, and the selection pressure rewarding a low computation and communication complexity would guide an evolution of adequate responses.

CONCLUSION

This chapter has presented an approach to the structural health management (SHM) of future aerospace vehicles that will need to operate robustly in very adverse environments. Such systems will need to be intelligent and to be capable of self-monitoring and ultimately, self-repair. The robustness requirement is best satisfied by using a distributed rather than a centralised system, and this has been assumed from the outset. Networks of embedded sensors, active elements, and intelligence have been selected to form a prototypical “smart skin” for the aerospace structure, and a methodology based on multi-agent networks developed for the system to implement aspects of SHM by processes of self-organisation. This has been developed in the context of a hardware test-bed, the CSIRO/NASA “concept demonstrator” (CD), a cylindrical structure with a metallic smart skin with 196 sensor/actuator/processor modules. A number of SHM algorithms related to damage detection and assessment have been developed and tested on this demonstrator.

A future aerospace vehicle will be expected to respond to a variety of damage situations which, moreover, vary with time and circumstance. Designing a general system with distributed intelligence which can self-organise solutions to many different problems is a very difficult task which we have simplified considerably by dividing the problems into manageable components as described in the *Response Matrix Approach* section, then seeking self-organising solutions to each component. This top-down/bottom-up (TDBU) approach allows solutions to be achieved whilst retaining the flexibility and emergent behaviour expected from complex multi-agent networks.

This breakdown of problems into components was achieved with the aid of a “response matrix” (Table 1) and three significant scenarios were analysed in this fashion. These were (a) critical damage, which threatens the integrity of the vehicle, (b) sub-critical damage, which requires immediate action although is not life-threatening, and (c) minor damage, whose cumulative effects need to be monitored and acted on when appropriate. From these scenarios, three main components were selected, and self-organising solutions developed for each and tested on the hardware test-bed. These components were: (1) the formation of “impact boundaries” around damage sites, allowing the extent of any damage to be assessed and communicated to other parts of the vehicle; (2) self-assembling “impact networks”, robust communications links which connect damage sites, enabling inspection of minor damage; and (3) shape replication, a demonstration of an autonomous repair mechanism by which the network “grows”, at a remote site, a new region of the correct shape to replace a damaged area. The first two of these have been successfully implemented on the hardware test-bed, giving confidence in the feasibility of the overall approach.

Future work will continue with the implementation of other necessary components, such as detailed diagnosis and prognosis for the different damage scenarios (Prokopenko et al., 2005b), sensor-data clustering (Mahendra, Prokopenko, Wang, & Price, 2005; Prokopenko, Mahendra, & Wang, 2005), robust communication to action-initiating sites (Li, Guo, & Poulton, 2004) and actions aimed at repair or mitigation of damage. One such development, currently in the preliminary stages, aims at developing a means of secondary inspection, an independent system which, when invoked by a report of possible damage, is capable of examining the relevant site and assessing the extent of the damage.

Looking further ahead, it is clear that the functionalities of sensing, computation, and action must merge with the material properties of the vehicle, moving closer to the real meaning of a smart skin. Although it may be some time before such a development is fully realised, recent progress in materials science and nano-technology gives confidence that it is achievable. We believe that the basic approach outlined in this chapter, of seeking self-organising solutions to critical components within an intelligent multi-agent framework, will still form the backbone of such future developments.

ACKNOWLEDGMENTS

A part of this work was carried out under CSIRO contract to NASA. It is a pleasure to record our appreciation to Dr. Ed Generazio (NASA Langley Research Center) for his encouragement and support. The authors are grateful to other members of the AAV project for their contributions to this work.

REFERENCES

- Abbott, D., Ables, J., Batten, A., Carpenter, D. C., Collings, A. F., Doyle, B., et al. (2003, January). *Development and evaluation of sensor concepts for ageless aerospace vehicles, Rep. 3: Design of the concept demonstrator* (Conf. Rep. No. TIPP 1628). CSIRO Telecommunications and Industrial Physics.
- Abbott, D., Doyle, B., Dunlop, J., Farmer, T., Hedley, M., Herrmann, J., et al. (2002). *Development and evaluation of sensor concepts for ageless aerospace vehicles: Development of concepts for an intelligent sensing system* (NASA Tech. Rep. No. NASA/CR-2002-211773). Hampton, VA: Langley Research Center.
- Batten, A., Dunlop, J., Edwards, G., Farmer, T., Gaffney, B., Hedley, M., et al. (2004, April). *Development and evaluation of sensor concepts for ageless aerospace vehicles, Rep. 5, Phase 2: Implementation of the concept demonstrator* (Conf. Rep. No. TIPP 2056). CSIRO Telecommunications and Industrial Physics.
- Buhl, J., Deneubourg, J. -L., & Theraulaz, G. (2002, September 12-14). Self-organized networks of galleries in the ant *messor sancta*. In M. Dorigo, G. Di Caro, & M. Sampels (Eds.), *Proceedings of the 3rd International Workshop on Ant Algorithms —ANTS2002*, Brussels, Belgium (pp. 163-175).
- Camazine, S., Deneubourg, J. -L., Franks, N. R., Sneyd, J., Theraulaz, G., & Bonabeau, E. (2001). *Self-organization in biological systems*. Princeton, NJ: Princeton University Press.
- Colorni, A., Dorigo, M., & Maniezzo, V. (1992). Distributed optimization by ant colonies. In F. S. Varela & P. Bourguine (Eds.), *Toward a Practice of Autonomous Systems: Proceedings of the 1st European Conference on Artificial Life*, Paris, December 11-13, 1991 (pp. 134-142).
- Crutchfield, J. (1994). The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75, 11-54.
- Deneubourg, J. -L., & Goss, S. (1989). Collective patterns and decision making. *Ethology, Ecology & Evolution*, 1, 295-311.

- Dorigo, M., & Di Caro, G. (1999, July 6-9). Ant colony optimization: A new meta-heuristic. *Proceedings of the 1999 Congress on Evolutionary Computation*, Washington, DC (pp. 1470-1477). Piscataway, NJ: IEEE Press.
- Dorigo, M., Maniezzo, M., & Coloni, A. (1996). The ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics, Part B*, 26(1), 1-13.
- Durbeck, L. J. K., & Macias, N. J. (2002, July 29-August 1). Defect-tolerant, fine-grained parallel testing of a cell matrix. In *Proceedings of SPIE, Series 4867—Convergence of IT and Communications (ITCom)*, Boston (pp. 71-85).
- Eppstein, D. (1996). *Spanning trees and spanners* (Tech. Rep. No. 96-16). Department of Information and Computer Science, University of California, Irvine. Also in J.-R. Sack & J. Urrutia (Eds.), *Handbook of computational geometry* (pp. 425-461). Amsterdam: Elsevier.
- Ferber, J. (1999). *Multi-agent systems*. Reading, MA: Addison Wesley.
- Foreman, M., Prokopenko, M., & Wang, P. (2003, September 14-17). Phase transitions in self-organising sensor networks. In W. Banzhaf, T. Christaller, P. Dittrich, J. T. Kim, & J. Ziegler (Eds.), *Advances in Artificial Life — Proceedings of the 7th European Conference on Artificial Life (ECAL-03)*, Dortmund, Germany (pp. 781-791). Heidelberg: Springer.
- Generazio, E. (1996). Nondestructive evaluation science in the year 2015+. *NASA Langley Research Center, Nondestructive Evaluation Sciences Branch*. Retrieved from <http://nesb.larc.nasa.gov/nde2015.html>
- Gulati, S., & Mackey, R. (2003, July). *Prognostics methodology for complex systems: Automatic method to detect and react to complex degradation and incipient faults* (Tech. Rep. No. NPO-20831). NASA Jet Propulsion Laboratory, Pasadena, California.
- Hedley, M., Hoschke, N., Johnson, M., Lewis, C., Murdoch, A., Price, D., et al. (2004, December 5-8). Sensor network for structural health Monitoring. In *Proceedings of the International Conference on Intelligent sensors, Sensor Networks and Information Processing (ISSNIP-2004)*, Melbourne, Australia (pp. 361-366).
- Hofstadter, D. R. (1989). *Gödel, Escher, Bach: An eternal golden braid*. New York: Vintage Books.
- Holland, O., & Melhuish, C. (1999). Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5, 173-202.
- Kahng, A. B., & Robins, G. (1995). *On optimal interconnections for VLSI*. Boston: Kluwer Academic Publishers.
- Klyubin, A. S., Polani, D., & Nehaniv, C. L. (2004, June 24-26). Organization of the information flow in the perception-action loop of evolved agents. *Proceedings of the NASA/DoD Conference on Evolvable Hardware*, Seattle, WA (pp. 177-180). Piscataway, NJ: IEEE Computer Society Press.
- Li, J., Guo, Y., & Poulton, G. (2004, December 4-6). Critical damage reporting in intelligent sensor networks. *Proceedings of the 17th Joint Conference on Artificial Intelligence (AI'04)*, Cairns, Australia (pp. 26-38).
- Lovatt, H., Poulton, G., Price, D., Prokopenko, M., Valencia, P., & Wang, P. (2003, July 14-18). Self-organising impact boundaries in ageless aerospace vehicles. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, & M. Yokoo (Eds.), *Proceedings of the*

- 2nd International Joint Conference on Autonomous Agents and Multi-Agent Systems*, Melbourne, Australia (pp. 249-256). New York: ACM Press.
- Mahendra Rajah, P., Prokopenko, M., Wang, P., & Price, D. (2005, September 14-16). Towards adaptive clustering in self-monitoring multi-agent networks. *Proceedings of the 9th International Conference on Knowledge Based and Intelligent Information and Engineering Systems (KES-2005)*, Melbourne, Australia.
- Mange, D., Sanchez, E., Stauffer, A., Tempesti, G., Marchal, P., & Piguet, C. (1998). Embryonics: A new methodology for designing field-programmable gate arrays with self-repair and self-replicating properties. *IEEE Transactions on VLSI Systems*, 6, 387-399.
- Price, D., Scott, A., Edwards, G., Batten, A., Farmer, A., Hedley, M., et al. (2003, September). An integrated health monitoring system for an ageless aerospace vehicle. In F.-K. Chang (Ed.), *Structural Health Monitoring 2003 — Proceedings of 4th International Workshop on Structural Health Monitoring*, Stanford, CA (pp. 310-318). Lancaster, PA: D. E. Stech Publications.
- Prokopenko, M., Mahendra, P., & Wang, P. (2005, September 5-9). On convergence of dynamic cluster formation in multi-agent networks. *Proceedings of the 8th European Conference on Artificial Life (ECAL 2005)*, Canterbury, Kent, UK.
- Prokopenko, M., & Wang, P. (2004, September 12-15). On self-referential shape replication in robust aerospace vehicles. *Proceedings of the 9th International Conference on the Simulation and Synthesis of Living Systems (ALIFE9)*, Boston (pp. 27-32).
- Prokopenko, M., Wang, P., & Price, D. (2005, June 29-July 1). Complexity metrics for self-monitoring impact sensing networks. In *Proceedings of the NASA/DoD International Conference on Evolvable Hardware (EH-05)*, Washington, DC.
- Prokopenko, M., Wang, P., Price, D., Valencia, P., Foreman, M., Farmer, A. (2005a). Self-organising hierarchies in sensor and communication networks. *Artificial Life (special issue on Dynamic Hierarchies)*, 11(4).
- Prokopenko, M., Wang, P., Scott, D. A., Gerasimov, V., Hoshke, N., & Price, D. C. (2005b, September 14-16). On self-organising diagnostics in impact sensing networks. *Proceedings of the 9th International Conference on Knowledge Based and Intelligent Information and Engineering Systems (KES-2005)*, Melbourne, Australia.
- Prosser, W. H., Allison, S. G., Woodard, S. E., Wincheski, R. A., Cooper, E. G., Price, D. C., et al. (2004, December 16-17). Structural health management for future aerospace vehicles. *Proceedings of the 2nd Australasian Workshop on Structural Health Monitoring (2AWSHM)*, Melbourne, Australia (pp. 1-15).
- Schoonderwoerd, R., Holland, O. E., Brutton, J., & Rothkrantz, L. J. M. (1997). Ant-based load balancing in telecommunication networks. *Adaptive Behaviour*, 5(2), 169-207.
- Schweitzer, F., & Tilch, B. (2002). Self-assembling of networks in an agent-based model. *Physical Review E*, 66, 1-9.
- Sipper, M., Mange, D., & Stauffer, A. (1997). Ontogenetic hardware. *BioSystems*, 44, 193-207.
- Trianni, V., Labella, T. H., Gross, R., Sahin, E., Dorigo, M., & Deneubourg, J. -L. (2002). *Modelling pattern formation in a swarm of self-assembling robots* (Tech. Rep. No. TR/IRIDIA/2002-12, May). Bruxelles, Belgium: Institut de Recherches Interdisciplinaires et de Developpement en Intelligence Artificielle.

- Venkatasubramanian, V., Katare, S., Patkar, P. R., & Mu, F. -P. (2004). Spontaneous emergence of complex optimal networks through evolutionary adaptation. *Computers and Chemical Engineering*, 28(9), 1789-1798.
- Wang, P., & Prokopenko, M. (2004, July 13-17). Evolvable recovery membranes in self-monitoring aerospace vehicles. In *Proceedings of the 8th International Conference on Simulation of Adaptive Behaviour (SAB-2004)*, Los Angeles, CA (pp. 509-518).
- Wang, P., Valencia, P., Prokopenko, M., Price, D., & Poulton, G. (2003, June 30-July 3). Self-reconfigurable sensor networks in ageless aerospace vehicles. In U. Nunes, A. T. de Almeida, A. K. Bejczy, K. Kosuge, & J. A. T. Machado (Eds.), *Proceedings of the 11th International Conference on Advanced Robotics (ICAR-03)*, Coimbra, Portugal (pp. 1098-1103).
- Wessnitzer, J., Adamatzky, A., & Melhuish, C. (2001, September 10-14). Towards self-organising structure formations: A decentralized approach. In J. Kelemen & P. Sosík (Eds.), *Advances in Artificial Life: Proceedings of the 6th European Conference on Artificial Life*, Prague, Czech Republic (pp. 573-581).
- Wu, Y. -F., Widmayer, P., Schlag, M. D. F., & Wong, C. K. (1987). Rectilinear shortest paths and minimum spanning trees in the presence of rectilinear obstacles. *IEEE Transactions on Computers*, C-36, 321-331.

Chapter VIII

Knowledge Through Evolution

Russell Beale, University of Birmingham, UK

Andy Pryke, University of Birmingham, UK

ABSTRACT

This chapter argues that a knowledge discovery system should be interactive, should utilise the best in artificial intelligence (AI), evolutionary, and statistical techniques in deriving results, but should be able to trade accuracy for understanding. Further, it needs to provide a means for users to indicate what exactly constitutes “interesting”, as well as understanding suggestions output by the computer. One such system is Haiku, which combines interactive 3D dynamic visualization and genetic algorithm techniques, and enables users to visually explore features and evaluate explanations generated by the system. Three case studies are described which illustrate the effectiveness of the Haiku system, these being Australian credit card data, Boston area housing data, and company telecommunications network call patterns. We conclude that a combination of intuitive and knowledge-driven exploration, together with conventional machine learning algorithms, offers a much richer environment, which in turn can lead to a deeper understanding of the domain under study.

INTRODUCTION

In this modern world, information is collected all the time: from our shopping habits to web browsing behaviours, from the calls between businesses to the medical records of individuals, data is acquired, stored, and gradually linked together. In this morass of data, there are many relationships that are not down to chance, but transforming data into information is not a trivial task. Data is obtained from observation and measurement, and has no intrinsic value. But from it we can create information: theories and relationships that describe the relationships between observations. And from information we can create knowledge: high-level descriptions of what and why, explaining and understanding the fundamental data observations. The mass of data available allows us to potentially discover important relationships between things, but the sheer volume dictates that we need to use the number-crunching power of computers to assist us with this process.

Data mining, or knowledge discovery as it is sometimes called, is the application of artificial intelligence and statistical analysis techniques to data in order to uncover information. Given a number of large datasets, we are fundamentally interested in finding and identifying interesting relationships between different items of data. This may be to identify purchasing patterns, which are then used for commercial gain through guiding effective promotions, or to identify links between environmental influences and medical problems, allowing better public health information and action. We may be trying to identify the effects of poverty, or to understand why radio-frequency observations of certain stars fluctuate regularly. Whatever the domain of the data, we are engaged in a search for knowledge, and are looking for interesting patterns in the data.

But what is “interesting”? One day, it may be that the data falls into a general trend; the next it may be the few outliers that are the fascinating ones. Interest, like beauty, is in the eye of the beholder. For this reason, we cannot leave the search for knowledge to computers alone. We have to be able to guide them as to what it is we are looking for, which areas to focus their phenomenal computing power on. In order for data mining to be generically useful to us, it must therefore have some way in which we can indicate what is interesting and what is not, and for that to be dynamic and changeable. Many data mining systems do not offer this flexibility in approach: they are one-shot systems, using their inbuilt techniques to theorise and analyse data, but they address it blindly, as they are unable to incorporate domain knowledge or insights into what is being looked for; they have only one perspective on what is interesting, and report only on data that fit such a view. Many such systems have been utilised effectively, but we believe that there is more to data mining than grabbing just the choicest, most obvious nuggets.

There are further issues with current approaches to data mining, in that the answers are often almost as incomprehensible as the raw data. It may be that rules can be found to classify data correctly into different categories, but if the rules to do so are pages long, then only the machine can do the classification: we may know how to do the classification, but have no insight into why it may be like that. We have gained information, but not knowledge. We believe that we should be able to understand the answers that the system gives us. In order to achieve this, it may be that we need broader, less accurate generalisations that are comprehensible to the human mind, but then feel confident in the main principles to allow the machine to do classification based on much more complex

rules that are refinements of these basic principles. For example, “if it’s red and squishy, it’s a strawberry” is easy to understand. Even if that’s true only 80% of the time, it’s a useful rule, and easier to grasp than:

```
red, deforms 4mm under 2N pressure, >3cm diameter = strawberry &
red, deforms 1mm under 2N pressure, <6cm diameter = cherry &
red, deforms 3 mm under 4N pressure, >5cm diameter = plum
else raspberry
```

which may be 96% correct but is hardly memorable. For many data mining systems, the rules developed are far more complex than this, each having numerous terms, with no overall picture able to emerge. For statistical-based systems, the parameter sets are even harder to interpret.

Since “interesting” is essentially a human construct, we argue that we need a human in the data mining loop; if we are to develop an effective system, we need to allow them to understand and interact with the system effectively. We should also take advantage of the capabilities of the user, many of which we have tried to emulate with AI systems for many years, and are still a long way from reproducing effectively. A key example is the human visual system, which is very effective at picking out trends within a mist of data points, capable of dealing with occlusion, missing values, and noise without conscious effort. On the other hand, processing vast numbers of points and deriving complex statistics is something much better suited to computers.

This leads us to conclude that a knowledge discovery system should be interactive, should utilise the best in artificial intelligence, evolutionary, and statistical techniques in deriving results, but should be able to trade accuracy for understanding; it also needs to provide a way of allowing the user to indicate what is interesting and to understand the suggestions that the computer makes. An ideal system should be symbiotic, each benefiting from the intrinsic abilities of the other, and holistic, producing results that are much more powerful than each could achieve on their own (Pryke & Beale, 2005).

KNOWLEDGE DISCOVERY WITH HAIKU

The Haiku system was developed with these principles in mind, and offers a symbiotic system that couples interactive 3D dynamic visualization technology with a novel genetic algorithm. The system creates a visualisation of the data which the user can then interact with, defining which areas are of interest and which can be ignored. The system then takes this input and processes the data using a variety of techniques, presenting the results as explanations to the user. These are in both textual and visual form, allowing the user to gain a broader perspective on what has been achieved. Using this information, they can refine what the system should look at, and slowly focus in on developing knowledge about whatever it is they are interested in. As well as using conventional rule generation techniques, Haiku also has a specifically designed genetic algorithmic approach to producing explanations of data. Each of these components is described in more detail as follows.

VISUALISATION

The visualisation engine used in the Haiku system provides an abstract 3D perspective of multi-dimensional data. The visualisation consists of nodes and links, whose properties are given by the parameters of the data. Data elements affect parameters such as node size, mass, link strength, elasticity, and so on. Multiple elements can affect one parameter, or a sub-set of parameters can be chosen.

Many forms of data can be visualised in Haiku. Typical data for data mining consists of a number of individual “items” (representing, for example, customers) each with the same number of numerical and/or nominal attributes. What is required for Haiku visualisation is that a distance can be calculated between any two items. The distance calculation should match an intuitive view of the differences between two items. In most cases, a simple and standard distance measure performs well: with data elements $\bar{x}_a = [x_1, x_2, \dots, x_n]$, the distance d between elements \bar{x}_a and \bar{x}_b is:

$$d = |\bar{x}_a - \bar{x}_b| = \sum_{i=1}^n x_{ai} - x_{bi} \quad (1)$$

An example of this is shown in Table 1.

The total distance $d = -26.53$. Clearly, many variations of this exist — a weighted sum can be used, and so on. One of the characteristics of the system is that the user can choose which parameters are used to create the distance metric, and which ones affect the other characteristics of the visualisation.

In the visualisation, a node is created that represents an item. These nodes may be all equivalent, or may have characteristics inherited from the data (for example, number of children may be used, not in the standard distance measure, but in the mass of the node). Links are created between all the nodes, which act as springs and try to move the nodes about in the space.

To create the visualisation, nodes are initially scattered randomly into the 3D space, with their associated links. This 3D space obeys a set of physical-type laws, which then affect this initial arrangement. Links tend to want to assume a particular length (directly related to the distance measure between the nodes), and tend to pull inwards until they reach that length, or push outwards if they are compressed, just as a spring does in the real world. Nodes tend to repel each other, based on their mass. This whole approach can be seen as a force-directed graph visualisation. This initial state is allowed to evolve, and the links and nodes shuffle themselves around until they reach a local minimum, low-energy steady state. The reasoning behind these choices of effects is that we want related

Table 1. Calculating distance between two data items

Data Item	Phone bill	Shopping	Petrol	Children	Age	Sum distance
Customer1	124.23	235.12	46.23	2	34	
Customer2	34.56	281.46	123.09	0	29	
Distance	89.67	46.34	76.86	2	5	219.87

things to be near to each other, and unrelated things to be far away. Therefore, by creating links that are attractive between data points with similar characteristics, we achieve this clumping effect. The data points themselves, the nodes in the visualisation, are made repulsive so that the system does not collapse to a point, but instead are individually distinguishable entities, slightly separated from their similar neighbours.

This approach achieves a number of things. It allows us to visualise high-dimensional data in a comprehensible and compact way. It produces results that are similar to those achieved using approaches such as multi-dimensional scaling, but is somewhat more comprehensible because it tries to cluster “similar” things with other “similar” ones. It is certainly true that the choice of distance metric, and particularly which items to include and which to map to node characteristics, can affect the resulting visualisation, but we are searching for insight and meaning, not trying to come up with a *single* right solution. At different times, different features can be examined, and different results achieved; this is an inherent characteristic of searching for information, rather than an intrinsic problem with the approach. In any move from a high-dimensional space to a lower one, information will have to be lost; this approach at least preserves some of the main similarity characteristics of the original datasets.

The physics of the space are adjustable, but are chosen so that a steady state solution can be reached that is static; this is unlike the real world, in which a steady state exists that involves motion, with one body orbiting another. This is achieved by working in a non-Newtonian space. In the real physical world (a Newtonian space), we have the following condition:

$$F = ma \quad (2)$$

where F is the force applied to a body, m the mass of that body, and a is the acceleration produced. This can be re-written as:

$$F = m \frac{dv}{dt} \quad (3)$$

where v is the velocity of the object.

When the visualisation is in a local minimum, there is no net force on any of the bodies (in other words, all the spring-like forces from the links and repulsive nodal forces balance each other out), and so for each node, $F = 0$. Thus:

$$0 = m \frac{dv}{dt} \Rightarrow \frac{dv}{dt} = 0 \Rightarrow v = \text{constant} \quad (4)$$

Therefore, in a steady-state Newtonian space, each node may potentially have zero or a constant velocity. In other words, the steady state solution has dynamic properties, with bodies moving in orbit, for example.

In our space, we redefine (2) to be:

$$F = mv \quad (5)$$

When we reach the steady state, we have (for non-zero masses):

$$0 = mv \Rightarrow v = 0 \quad (6)$$


Thus, in our representations the steady state that the arrangement evolves to is static.

This representation can then be explored at will by rotating it, zooming in, and flying through and around it. It is a completely abstract representation of the data, and so has no built-in preconceptions. Different data-to-attribute mappings will clearly give different structures, but the system can at least produce a view of more than three dimensions of the raw data at once. A typical structure is shown in Figure 1.

To evolve the structure, each node is checked for links to other nodes, and the forces of those links is added vectorially to give a net force, and the node is then moved according to that force using Equation 5 above. Computationally, the process scales exponentially with the number of links, which is usually proportional to the number of data points, so the evolution to the stable structure moves from being a real-time process that you can watch towards one that has to be allowed to run for a long period of time as the dataset increases in size.

In general, this is not a problem, since the initial arrangement of data is random and the evolutionary process is not in itself informative (although it is interesting to observe). However, when the visualisation is used as a component in the data mining tool, this is designed to be an interactive process, so we have taken a number of approaches to speeding up the relaxation to steady state. The first involves re-coding the system into OpenGL/DirectX, to take advantage of the power of modern graphics processors, especially for 3D work. The second places the nodes into the space in a non-random position initially; each node is placed “near” a node it has a link to. This is marginally more computationally expensive initially, but reduces the numbers of nodes that have to move a large amount through the visualisation, and hence cause large scale changes in other nodal positions. The most effective approach is to use predominantly local

Figure 1. Nodes and links self-organised into a stable structure



relaxation; however, instead of considering all the forces to act over infinite distance, we can limit nodal interactions to be very local, so that nodes which are a long way away do not exert any forces on the ones in question (much like assuming that the gravitational effects of all the stars except the sun are negligible).

Once the system has undergone some initial relaxation, which provides some level of organisation, we can also focus on the local neighbourhood much more, and occasionally recompute the longer-range interactions. This is akin to organising a tight cluster properly, but then treating that as one structure for longer-range effects. A combination of these approaches allows us to produce an effective steady state representation, even with large datasets, in interactive time.

PERCEPTION-ORIENTED VISUALISATION

The interface provides full 3D control of the structure, from zooming in and out, moving smoothly through the system (flyby), rotating it in 3D, and jumping to specific points, all controlled with the mouse. Some typical structures emerge, recognisable from dataset to dataset. For example, a common one is the “dandelion head”: a single central node connected to a number of other nodes with the same strength links. The links pull the attached nodes towards the central one, but each node repels the others, and so they spread out on the surface of a sphere centred on the main node. This looks much like a dandelion head. Another typical structure occurs when a number of dandelion heads are loosely linked together. The effect of the other heads in the chain forces the outer nodes away from being equidistantly spaced on the sphere and makes them cluster together somewhat on the side away from the link, and a series of “florets” are created, all linked together. It is because of this that some users have termed the visualisation “cauliflower space”.

The visualisation in itself provides a lot of information about the dataset. We have used the visualisation in isolation for a number of tasks (Hendley, Drew, Beale, & Wood, 1999). One of the more effective ones has been the visualisation of users’ internet browsing behaviour. Each page visited is represented by a node, and their page transitions are represented by the links. Typically, users start on a home or an index page, and move out and back a number of times before moving off down a promising thread: this behaviour, when visualised in real time, produces a dandelion head with increasing numbers of “seeds” (the outer nodes) and then switches towards a floret as the thread is followed. A new index-type page is reached (sometimes after one hop, sometimes after many, and another floret is created. Often, there are links back to the originally explored pages, and when the user follows these, the visualisation pulls itself into a ring, representing a notion of closure and returning that has an exact analogy in the real world (Wood, Drew, Beale, & Hendley, 1995). A different representation is formed if we visualise the structure of web pages: pages themselves are nodes again, but hyperlinks map to visualisation links. A Web site has a fairly typical cauliflower image, caused by closely interrelated and interlinked sections, tied back to a common home or index page, with links off to other cauliflowers where the site links externally to other sites.

The system has also been used to assist users to comprehend their progress in information retrieval tasks. Using a digital library as our domain, for each query a representation of the results was returned. A large node represented the query, and was

Figure 2. Visualising the result of a single query: “visualisation colour graphics”




Figure 3. Adding a second query: “3D surface graphics”




Figure 4. Adding a third, unrelated query: “agents”





Figure 5. A sequence of four queries



fixed in the 3D space. Each document that matched the query was a mobile node, with a link attaching it to the query, with the link strength being how relevant the document was to that query. An initial query would return a number of documents, so a distorted dandelion head would appear. However, a second query that returned some of the same documents would show links from those documents to both fixed nodes, and hence the degree of overlap could be easily seen. Such an approach allowed the user, in real time, to see how effectively they were exploring the space of documents, and how those were interrelated to the queries made (Beale, McNab, & Witten, 1997; Cunningham, Holmes, Littin, Beale, & Witten, 1998). This is important as subsequent searches are often dependent on the results of the previous ones, so having a representation of the history and its relationships to the present search matches more closely what the user is doing internally. A walkthrough of the process is shown in Figures 2 through 5.

Interaction with the Data Visualisation

When features of interest are seen in the visual representation of the data, they can be selected using the mouse. This opens up a number of possibilities:

- Data identification
- Revisualisation
- Explanation

The simplest of these (data identification) is to view the identifiers or details or items in the feature, or export this information to a file for later use.

Another option is re-visualise the dataset without the selected data or to focus in and only visualise the selected data. This can be used to exclude distorting outliers, or to concentrate on the interactions within an area of interest. Of course, we can data mine the whole dataset without doing this, the approach taken by many other systems. One of the features of the Haiku system is this interactive indication of the things that we are currently interested in, and the subsequent focussing of the knowledge discovery process on best describing that data only.

A key feature of the system is that this user selection process takes full advantage of the abilities of our visual system: humans are exceptionally good at picking up gross features of visual representations. Our abilities have evolved to work well in the presence of noise, of missing or obscured data, and we are able to pick out simple lines and curves, as well as more complex features such as spirals and undulating waves or planes. By allowing user input into the knowledge discovery process, we can effectively use a highly efficient system very quickly as well as reducing the work that the computational system has to perform.

Explanation

The most striking feature of the system is its ability to “explain” why features of interest exist. Typical questions when looking at a visual representation of data are: “Why are these items out on their own?”, “What are the characteristics of this cluster?”, “How do these two groups of items differ?”. Answers to these types of question are generated by applying a machine learning component.

The interaction works as follows: First, a group or number of groups is selected. Then the option to explain the groups is selected. The user answers a small number of

questions about their preferences for the explanation (short/long) (highly accurate/characteristic), and so on. The system returns a set of rules describing the features selected.

As an alternative, the classic machine learning system C4.5 (Quinlan, 1992) may be used to generate classification rules. Other data mining systems may be applied by saving the selected feature information to a comma-separated value file.

Rule Visualisation

Rules generated using C4.5 or the GA-based method can be visualised within the system to give extra insight into their relationships with the data. Rules are usually represented by massive nodes that do not move far in space, and are regularly spaced. Links show which rules apply to which data, and hence unclassified data and multiply-classified data are shown well.

From this, the processing moves towards the computer, as the genetic algorithm-based process takes over.

GENETIC ALGORITHMS FOR DATA MINING

We use a genetic algorithm (GA) approach for a number of reasons. The first is that a GA is able to effectively explore a large search space, and modern computing power means we can take advantage of this within a reasonable timeframe. We use a special type of GA that evolves rules; these produce terms to describe the underlying data of the form:

IF term OP value|range (AND ...) THEN term OP value|range (AND ...) (7)

where term is a class from the dataset, OP is one of the standard comparison operators ($<$, $>$, $=$, \leq , \geq), value is a numeric or symbolic value, and range is a numeric range. A typical rule would therefore be:

IF colour = red AND consistency = soft THEN fruit = strawberry (8)

A set of these rules can, in principle, describe any arbitrary situation. There are two situations that are of interest to us; classification, when the left hand side of the equation tries to predict a single class (usually known) on the right hand side, and association, or clustering, when the system tries to find rules that characterise portions of the dataset. The algorithm follows fairly typical genetic algorithmic approaches in its implementation, but with specialised mutation and crossover operators, in order to explore the space effectively. We start with a number of random rules, and evolve the population through subsequent generations based on how well they perform.

The genetic algorithm aims to optimise an objective function, and manipulation of this function allows us to explore different areas of the search space. For example, we can strongly penalise rules that give false positive results and achieve a different type of description, than rules that may be more general and have greater coverage but make a few more mistakes. Each rule is analysed in terms of the objective function and given a score, which is its *fitness*. The fittest rules are then taken as the basis for the next population, and new rules are created. Crossover points are chosen to be in syntactically-

similar positions, in order to ensure that we are working with semantically-meaningful chunks. Mutation is specialised: for ranges of values, it can expand or contract that range; for numbers, it can increase or decrease them; for operators, it can substitute them with others.

Statistically principled comparisons showed that this technique is at least as good as conventional machine learning at classification (Pryke, 1999), but has advantages over the more conventional approaches in that it can perform clustering operations as well. One of the key design features is to produce a system that has humanly-comprehensible results. Rules of the form in Equation 7 are inherently much more understandable than decision trees or probabilistic or statistical descriptions. It is also true that short rules are going to be easier to comprehend than longer ones. Since the GA is trying to minimise an objective function, we can manipulate this function to achieve different results. If we insist that the rules produced must be short (and hence easier to understand), then the system will trade off accuracy and/or coverage but will give us short rules, because they are “fitter”, which provide a general overview that is appropriate for much of the data. Because the Haiku system is interactive and iterative, when we have this higher level of comprehension, we can go back into the system and allow the rules to become longer and hence more specific, and accuracy will then increase.

FEEDBACK

The results from the GA are fed back into the visualisation: identified clusters can be coloured, for example, or rules added and linked to the data that they classify, as in Figure 6.


In this figure, rules are the large purple (left and centre), fuschia (rightmost) and green (top) spheres, with the data being the smaller spheres. Links are formed between the rules and the data that is covered by the rule, and the visualisation has reorganised itself to show this clearly. We have additionally coloured the data according to its correct classification (clear in colour, harder to see in greyscale).

A number of things are immediately apparent from this visualisation, much more easily than would be the case from a textual description. On the very left of Figure 6, one rule covers exactly the same data as the second sphere from the left, except it also misclassifies one green data point. But this second sphere, while correctly classifying all its own data correctly, also misclassifies much of the other data as well, shown by the many links to the different coloured data items. The visualisation shows us that we can remove this rule, simplifying the description, without reducing coverage and improving accuracy. On the right hand side of the picture, the rule clearly does very well; it covers all its data and does not misclassify anything. The rule at the top has mixed results.

The system is fully interactive, in that the user can now identify different characteristics and instruct the GA to describe them, and so the process continues.

This synergy of abilities between the rapid, parallel exploration of the structure space by the computer and the user’s innate pattern recognition abilities and interest in different aspects of the data produces a very powerful and flexible system.

Figure 6. Rules and classified data



CLASSIC CASE STUDY 1: WELL KNOWN DATASETS

Several machine learning datasets from the UCI Machine Learning Repository (Blake & Merz, 1998) were used to benchmark the performance of data mining and classification. It should be noted that it focuses on quantitative performance, whereas the qualitative experience and use of perception-based mining techniques is not assessed. However, good results on these datasets in quantitative terms will give us confidence when analysing new datasets.

The GA-based approach gave perfectly acceptable results, with statistical analysis showing it performed better than C4.5 (Quinlan, 1992) on the “Australian Credit Data” ($p=0.0018$). No significant difference in performance was found for the other two datasets. These results are summarised in Table 2.

CASE STUDY 2: INTERACTIVE DATA MINING OF HOUSING DATA

Figure 7 shows a 2D view of the system’s visual clustering of the Boston housing data. Two user selected groups have been indicated.

GA-based data mining was then applied to these user identified groups. The fitness function was chosen so as to bias the system towards the discovery of rules which are short and accurate (Table 3).

This case study illustrates the following qualitative aspects of the system. The interactive visual discovery approach has revealed new structure in the data by visual clustering. Subsequent application of the data mining algorithm has generated concrete information about these “soft” discoveries. These rules look at a variety of aspects of

Table 2. Quantitative benchmarking performance

Dataset	Genetic algorithm % correct	C4.5 % correct
Australian Credit (Quinlan, 1987)	86%	82%
Boston Housing (Quinlan, 1993)	64%	65%
Pima Indians Diabetes (Smith et al., 1988)	73%	73%

the system, from their location to their tax rates to their social status, and provide rules that are accurate, short, and cover much of the data, and they are comprehensible. Together, interactive data mining has delivered increased knowledge about a well-known dataset.

Having proven its worth on known datasets, we have used the system to try to discover *new* phenomena.

CASE STUDY 3: APPLYING HAIKU TO TELECOMS DATA

Justification

Massive amounts of data are generated from monitoring telecommunications switching. Even a small company may make many thousands of phone calls during a year. Telecommunications companies have a mountain of data originally collected for billing

Figure 7. Clusters selected in the Boston housing data




Table 3. Rules generated from Boston housing data

Rule	Accuracy (%)	Coverage (%)
Bounds_river=true \Rightarrow GROUP_1	100	43
PropLargeDevelop = 0.0 AND 9.9 \leq older_properties_percent \leq 100.0 AND Pupil_teacher_ratio = 20.2 \Rightarrow GROUP_1	94	83
Bounds_river=false AND 4 \leq Highway_access \leq 8 \Rightarrow GROUP_2	100	77
Bounds_river=false AND 264 \leq Tax_rate \leq 403 \Rightarrow GROUP_2	100	69
2.02 < Industry_proportion \leq 3.41 \Rightarrow GROUP_2	98	13
5.68 \leq Lower_status_percent \leq 6.56 \Rightarrow GROUP_2	96	75
Bounds_river=false \Rightarrow GROUP_2	73	100

purposes. Telecoms data reflects business behaviour, so is likely to contain complex patterns. For this reason, Haiku was applied to mine this data mountain.

Data

The data considered detailed the calling number, recipient number, and duration of phone calls to and from businesses in a medium sized town. Other information available included business sector and sales channels. All identity data was anonymized.

Call Patterns of High Usage Companies

Visualisation

A number of companies with particularly high numbers of calls were identified. These were visualised separately to identify patterns within the calls of individual companies.

Figure 8 shows a clustering of calls from a single company. The most immediately obvious feature is the “blue wave” to the right of the image; this has been labelled “A”.


Also visible are various other structures, including the two clusters labelled “B” and “C”.

Discoveries

After identifying these features, we then asked the system to “explain” their characteristics. The following rules were discovered by the system, and translated into sentence form for clarity:

- All calls in group A are to directory enquiries.
 - Further investigation, selecting parts of the “blue wave” showed that the wave structure was arranged by hour of day in one dimension and day of week in the other.

Figure 8. Calls from one company, automatically clustered by Haiku (three areas are apparent — labelled A, B and C)



- Within group B, about 70% of calls are to two numbers. 90% of calls to these numbers fall into the group B. Almost all of the remaining 30% of calls in group B are to another two numbers.
- Most long-distance ISDN calls are in group B. All but one call in the group has these properties. Most calls in the group are also charged at the same rate.
- About 80% of Group C calls are ISDN calls, and about 10% are from payphones. About one third occur between 21:00 and 22:59, and about one half start at 15 minutes past the hour. Most are long-distance calls. About 50% of the calls are very long, lasting between 8 and 15.5 hours.

For this dataset, Haiku discovers some very interesting facts about the calling patterns of a company. Notice that we can produce short, comprehensible rules that cover a significant portion of the dataset, which are intrinsically much more usable than detailed descriptions of 100% of the data. These insights can then be used by the company to optimise their phone usage, or, as for this study, to feed back to the telecoms company some concepts for marketing and billing strategies.

CONCLUSION

The Haiku system for information visualisation and explanation provides a useful interface for interactive data mining. By interacting with a virtual data space created dynamically from the data properties, greater insight can be gained than by using standard machine learning- based data mining. It allows users to explore features visually, to direct the computer to generate explanations and to evaluate the results of

their exploration, again in the visual domain. By using a novel genetic algorithmic approach, we can bias rules generated to give us first a general overview and then progressively refine their accuracy and coverage as our understanding increases. This combination of intuitive and knowledge-driven exploration with the mechanical power of the learning algorithms provides a much richer environment and can lead to a deeper understanding of the domain.

ACKNOWLEDGMENTS

This work was partially supported by grants from British Telecom, Integral Solutions, Ltd., and British Maritime Technology. Thanks to Nick Drew and Bob Hendley for their work on the visualisation parts of the system, and to colleagues for their comments and help.

REFERENCES

- Beale, R., McNab, R. J., & Witten, I. H. (1997). Visualising sequences of queries: A new tool for information retrieval. *Proceedings of the IEEE Conference on Information Visualisation*, London, August (pp. 57-62).
- Blake, C. L., & Merz, C. J. (1998). UCI repository of machine learning databases. *University of California, Irvine, Department of Computer Science*. Retrieved from <http://www.ics.uci.edu/~mlearn/MLRepository.html>
- Cunningham, S. J., Holmes, G., Littin, J., Beale, R., & Witten, I. H. (1998). Applying connectionist models to information retrieval. In S. -I. Amari & N. K. Kasabov (Eds.), *Brain-like computing and intelligent systems* (pp. 435-457). Berlin: Springer-Verlag.
- Hendley, R. J., Drew, N., Beale, R., & Wood, A. M. (1999). Narcissus: Visualising information. In S. Card, J. Mackinlay, & B. Shneiderman (Eds.), *Readings in information visualization* (pp. 503-511). San Mateo, CA: Morgan Kaufmann.
- Pryke, A. N. (1999). *Data mining using genetic algorithms and interactive visualisation*. PhD thesis, Department of Computer Science, University of Birmingham, UK.
- Pryke, A. N., & Beale, R. (2005). Interactive comprehensible data mining. In C. Yang (Ed.), *Ambient intelligence for scientific discovery: Foundations, theories, and systems* (pp. 48-65). Berlin: Springer-Verlag.
- Quinlan, J. R. (1987). Simplifying decision trees. *International Journal of Man-Machine Studies*, 27, 221-234.
- Quinlan, J. R. (1992). *C4.5: Programs for machine learning*. San Mateo, CA: Morgan Kaufmann.
- Quinlan, J. R. (1993). Combining instance-based and model-based learning. *Proceedings of the 10th International Conference of Machine Learning, University of Massachusetts, Amherst* (pp. 236-243). San Mateo, CA: Morgan Kaufmann.
- Smith, J. W., Everhart, J. E., Dickson, W. C., Knowler, W. C., & Johannes, R. S. (1988). Using the ADAP learning algorithm to forecast the onset of diabetes mellitus. *Proceedings of the Symposium on Computer Applications and Medical Care* (pp. 261-265). Piscataway, NJ: IEEE Computer Society Press.

Wood, A. M., Drew, N. S., Beale, R., & Hendley, R. J. (1995). HyperSpace: Web browsing with visualisation. *Proceedings of the 3rd International World-Wide Web Conference*, Darmstadt, Germany, April (pp. 21-25).

Chapter IX

Neural Networks for the Classification of Benign and Malignant Patterns in Digital Mammograms

Brijesh Verma, Central Queensland University, Australia

Rinku Panchal, Central Queensland University, Australia

ABSTRACT

This chapter presents neural network-based techniques for the classification of micro-calcification patterns in digital mammograms. Artificial neural network (ANN) applications in digital mammography are mainly focused on feature extraction, feature selection, and classification of micro-calcification patterns into 'benign' and 'malignant'. An extensive review of neural network based techniques in digital mammography is presented. Recent developments such as auto-associators and evolutionary neural networks for feature extraction and selection are presented. Experimental results using ANN techniques on a benchmark database are described and analysed. Finally, a comparison of various neural network-based techniques is presented.

INTRODUCTION

Every year many women die from breast cancer worldwide. A recent study on breast cancer shows that one of every three cancer diagnoses in women is a breast cancer (<http://www.breastcancerfund.org> - Breast Cancer Facts 2003). Reports by various cancer institutes estimate that one in eight women develops breast cancer in the U.S. (http://www.breastcancerfund.org/disease_facts.htm - Breast Cancer Facts 2002), one in nine women in the UK and Canada (<http://www.cancerscreening.nhs.uk/breastscreen/breastcancer.html>), and one in ten women in Australia (<http://www.nbcc.org.au/>). The Australian National Breast Cancer Centre also reports that nearly 3% of women die from breast cancer worldwide, with the risk increasing with age, particularly after 50.

Digital mammography is considered to be one of the most reliable methods for early detection of breast cancer. The introduction of mammography screening in 1963 brought a major revolution to breast cancer detection and diagnosis. It has been widely adopted in many countries, including Australia, as a nationwide public health care program. According to the American College of Radiology, the decline in the number of breast cancer deaths corresponds *directly* to an increase in routine mammography screening (<http://www.acr.org>).

In digital mammography, most breast cancers are detected by the presence of micro-calcifications, which are one of the mammographic hallmarks of early breast cancer; they appear as a small bright spot on the mammogram. To decide whether a suspicious area on a digital mammogram contains a benign or malignant breast abnormality, traditionally the tissue has to be removed for examination using breast biopsy techniques.

Advanced image processing techniques are able to detect breast abnormalities efficiently; though their classification as “malignant” or “benign” still remains a challenging problem (Aghdasi, Ward, & Palcic, 1994; Gonzalez & Woods, 1993; Jain, 1995; Karssemeijer, 1994; Karssemeijer, Thijssen, Hendriks, & van Erning, 1998; Kopans, 1998; Lee & Bottema, 2000; Masek, Attikouzel, & deSilva, 2000; Neiber, Mueller, & Stotzka, 2000; Pereira & Azevedo, 2000; Sonka & Fitzpatrick, 2000; Umbagh, 1998; Wei, Laurence, & Clark, 1994; Yin, Giger, Vyborny, Doi, & Schmidt, 1993; Yoon, Ro, Kim, & Park, 2002; Zheng, Qian, & Clarke, 1994).

The abundance of variety and lack of individuality in micro-calcification patterns make their classification challenging for expert radiologists, even in high-resolution mammograms. Worldwide mass usage of screening mammography generates numerous amounts of mammograms every year, which requires a large number of skilled radiologists for interpretation. The variety of abnormal structures, long reading time, and monotony of interpretation work often produces human errors, missing either malignant cases or more benign biopsies. Therefore, there is a critical need for an intelligent system which can interpret mammograms accurately and uniformly using expert knowledge based on learning from experience. Along with expert radiologists, a computer-aided intelligent classification technique can be effectively used to improve and speed up the overall interpretation process.

Artificial neural networks (ANNs) have extraordinary generalization capabilities, which make them very suitable for use in computer-aided intelligent systems for breast cancer diagnosis (Bakic & Barzakovic, 1997; Cheng, Cai, Chen, Hu, & Lou, 2003; Wei, Nishikawa, & Doi, 1996; Wu, 1993). ANNs are adaptive intelligent tools that learn from examples (training set) and generalize new cases (test set) which they have never seen

before. Recently artificial neural networks have been used in the detection and classification of calcification and mass types of breast abnormalities. Some ANNs, along with other intelligent techniques, produce promising results in distinguishing benign from malignant patterns.

The remainder of this chapter is divided into four sections. The first presents a general overview of research methodology using neural networks. This section presents various benchmark databases, detection algorithms, feature extraction, selection, and classification techniques. The next section presents some recent results using neural evolutionary and auto-associator-based neural classification techniques. A comparative analysis of neural techniques is presented in the following section. The final section concludes the chapter.

RESEARCH METHODOLOGY USING ARTIFICIAL NEURAL NETWORKS

A general overview of research methodology for the classification of micro-calcification patterns using neural networks is presented in Figure 1 and described below. The research methodology comprises the following stages:

1. Digital mammography database.
2. Detection algorithm, or “suspicious area” marked by radiologists.
3. Area extraction using detection algorithm or chain code provided by radiologists.
4. Feature extraction (*optional*).
5. Feature selection (*optional*).
6. Classification of patterns into “benign” and “malignant”.

Figure 1. Research methodology overview

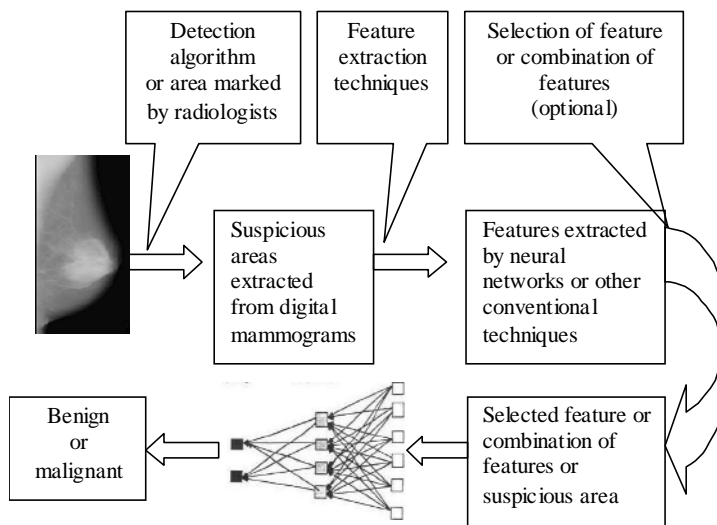



Figure 2. Digital mammograms (L-R: normal; architectural distortion; calcifications)




Digital Mammography Database

Digital Mammogram

Digital mammograms are created by using a camera, scanner, or specific mammogram digitizer. Usually 12 bits (4096 grey levels) per pixel are used to produce a high-resolution digital mammogram without loss of information. Figure 2 shows digital mammograms containing normal, architectural distortion, and calcifications patterns. The mammograms are taken from the Mammographic Image Analysis Society benchmark database (<http://www.wiau.man.ac.uk/services/MIAS/MIASweb.html>).

Digital mammograms are very useful in identifying breast abnormalities earlier than they can be diagnosed by physical examination. As shown in Figure 3, breast abnormalities can be divided into three types, namely: micro-calcification, circumscribed lesions, and spiculated lesions. Most breast cancers are detected by the presence of micro-calcifications.

Figure 3. Breast abnormalities (L-R: micro-calcification; circumscribed lesion; spiculated lesion)



Benchmark Databases

The implementation and testing of new algorithms/techniques, including artificial neural networks for the detection and classification of benign and malignant patterns, requires an appropriate training and testing database. Various international organizations and societies have created several digital mammography databases; some of the most popular of these are described below.

The Digital Database for Screening Mammography (DDSM) is one of the most popular digital mammography databases for training and testing new breast cancer detection techniques (Heath, Bowyer, Kopans, Moore, & Kegelmeyer, 2000). The database is available free of cost from the following website: <http://marathon.csee.usf.edu/Mammography/Database.html>

DDSM has already been divided into training and test mammograms, which makes it easier to test and compare various techniques. The database contains approximately 2,500 studies. Each study includes two images (MLO — medio-lateral oblique and CC — cranio caudal views) of each breast, along with some associated patient information such as age at time of study, density rating, subtlety rating for abnormalities, description of abnormalities, image information, type of scanner, spatial resolution, and so forth. Images containing suspicious areas have associated pixel-level “ground truth” information about the locations and types of suspicious regions.

The Nijmegen database has been a very popular database for a number of years; however, currently it is unavailable. It was removed from the web on March 31, 2000, because much better databases — such as DDSM — had become available. The DDSM database provides a much larger, and thus more diverse, set of mammograms with ground truth.

The UCI Repository contains two breast cancer databases (the Breast Cancer Database from the Ljubljana Oncology Institute, and the Wisconsin Breast Cancer Database), which are mainly used by the machine learning community for the analysis of ML algorithms. The database contains nine input feature values and two output values (“benign” and “malignant” classes). These databases do not contain digital mammograms, so they cannot be used to test new feature extraction techniques or detection algorithms. A detailed description of this database can be found in Blake and Merz (1998).

The Mammographic Image Analysis Society (MIAS) Database is an organization of UK research groups interested in the understanding of digital mammograms. MIAS has generated a database of digital mammograms (<http://www.wiau.man.ac.uk/services/MIAS/MIASweb.html>). Films taken from the UK National Breast Screening Program have been digitized to 50 micron pixel edges using a Joyce-Loebl scanning micro-densitometer — a device linear in the optical density range 0-3.2 and representing each pixel with an 8-bit word. The database contains 320 digitized films and is available on a DAT-DDS tape. It also includes radiologist “truth” markings on the locations of any abnormalities that may be present.

Detection of Breast Abnormalities

The aim of manual or automatic detection of breast abnormalities is to find suspicious areas on a mammogram which may or may not contain a malignant pattern. Computer-based detection is done by using various statistical and intelligent techniques, including artificial neural networks. Manual detection is done by two to three

expert radiologists based on their knowledge and experience. The program or expert radiologists provide centre and radius or chain code for detected breast abnormalities. The benchmark databases described above (DDSM, MIAS) contain suspicious areas marked by expert radiologists.

ANNs have been used for the development of detection algorithms. Wei, Nishikawa, and Doi (1996) used a shift invariant feed-forward network for the detection of clustered micro-calcifications. They obtained over an 88% detection rate; however, the algorithm was only tested on a small benchmark database. Sajda, Spence, and Pearson (1995) and Sajda, Spence, Pearson, and Nishikawa, (1996) used hierarchical feed-forward networks and contextual information to improve detection. Diahi, Frouge, Giron, and Fertil, (1996) also used hierarchical networks and obtained over 90% detection rate for micro-calcifications and masses. Comparative analyses of various neural networks and other detection algorithms can be found in Woods, Doss, Bowyer, Solka, Priebe, and Kegelmeyer, (1993), Bakic and Barzakovic (1997), and Cheng et al. (2003).

Christoyianni, Dermatas, and Kokkinakis (1999; 2000) used radial basis function neural networks with a set of decision criteria for detecting circumscribed masses in mammograms. They reported a 72.72% detection rate on 22 mammograms containing circumscribed lesions taken from the MIAS database. The lesion sizes varied from 18 to 198 pixels in radius. Kim and Park (1999) used back-propagation (BP) neural networks to evaluate the performance of different types of feature extraction techniques in their study of clustered micro-calcification detection in digitized mammograms. They have attained the highest (0.93) A_z value (conventional bi-normal model used as an index of diagnostic performance for the receiver operating characteristic (ROC) curve) with features extracted through the surrounding region-dependence method.

Yu and Guan (2000) used general regression neural networks to analyse the discriminatory power of wavelet and statistical features via sequential forward/backward selection methods for detecting micro-calcifications. Their method attained a 90% true positive detection rate at the cost of 0.5 false positive per image, while using 40 images containing 105 clusters of micro-calcifications from the Nijmegen database.

Comparative analyses of various neural networks and other detection algorithms can be found in Woods et al. (1993), Bakic and Brzakovic (1997), and Cheng et al. (2003).

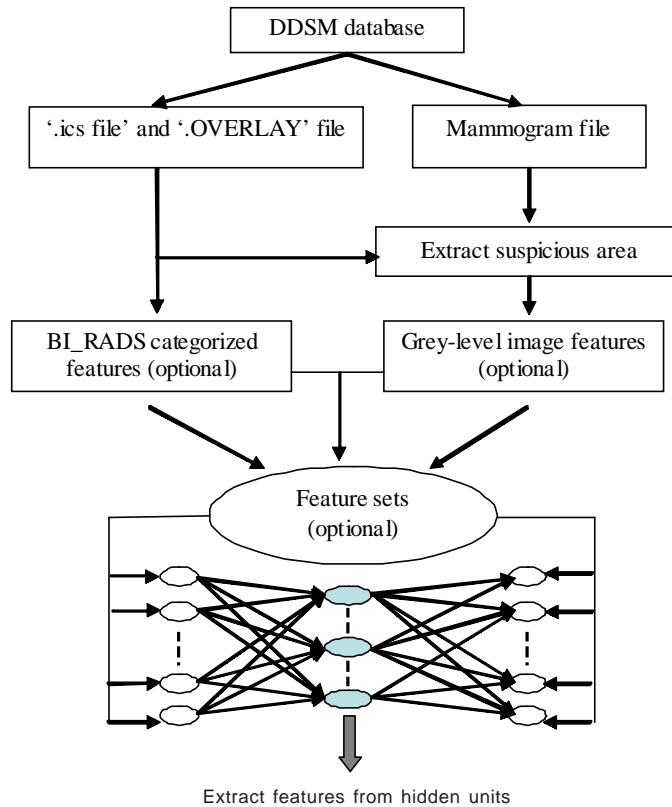
Area Extraction Based on Chain Code from the DDSM Database

The DDSM database has an .OVERLAY file for every mammogram, which contains information about the chain code used to extract the exact area. With the help of chain code values, the boundary of each suspicious area of the mammogram can be readily located.

Neural Network as Feature Extractor

The feature extraction process is optional in classifying “benign” and “malignant” patterns; however, it is one of the most important processes. The role of the feature extraction process in digital mammography is to take suspicious areas identified by the detection algorithms or expert radiologists and extract some important features (single or multiple values) which may help classifiers to distinguish malignant patterns from

Figure 4. Feature extraction using an auto-associator



normal (benign) patterns. Mainly, statistical feature extraction techniques have been used in the past few decades. The most common features used in digital mammography are entropy, contrast, moment, correlation, mean, deviation, area, shape, smoothness, average grey-level, texture, and BI-RADS (breast imaging reporting and data system) features.

Neural networks have recently been applied as feature extractors in digital mammography (Cheng et al., 2003; Panchal & Verma, 2005). Self-organizing maps (SOMs) and auto-associators are the most common types of neural networks used for feature extraction in digital mammography. SOM is based on unsupervised learning, in which an input image is taken and mapped into a small output vector (Cheng et al., 2003); a simple Kohonen learning algorithm is used. Auto-associators are based on supervised learning, and can compress and extract structural features from the input data (Panchal & Verma, 2005). In an auto-associator, the same inputs and target outputs are used to train the network. Back-propagation or RBF-based training is used to minimize the error between inputs and targets. Once we reach an acceptable error, the training is stopped and the

features are extracted from the hidden units. More details of using an auto-associator to extract features from a digital mammogram (from the DDSM database) are provided in Figure 4.

Neural Networks for Feature Selection

Feature selection is an important process in any classification problem because some irrelevant features may affect the classification rate and increase the cost. The feature selection process requires careful selection of which features to use to represent the patterns to be classified. Feature selection aims to find the best feature or combination of features which may achieve the highest classification rate. Many statistical techniques, such as best-first search, feature weighting, and hill-climbing, have been used for feature selection in digital mammography.

There has been also some research using neural techniques in the area of digital mammography to find the best feature or combination of features which can improve the classification rate. Artificial neural networks have been used for feature selection as a stand-alone technique (Verma & Zakos, 2001), or in conjunction with evolutionary algorithms (Zhang, Verma, & Kumar, 2004).

The following steps can be used for feature selection using a neural evolutionary technique:

1. Initialize the population.
2. Generate the inputs and get the parameters from the trained ANNs.
3. Compute the fitness for every individual of the current population.
4. Reproduce the new population by performing selection strategy, crossover, and mutation operators.
5. Calculate the fitness of the new population.
6. Check the number of generation and go back to step 2.

Neural Network as Classifier

Neural network classifiers have been widely applied in real-world applications. The purpose of this section is to review neural classifiers in digital mammography. ANNs have been used in classifying suspicious areas using either the whole area or just features.

Chitre, Dhawan, and Moskowitz (1993) used a back-propagation neural network for image structure features for micro-calcification classification, and compared their results with statistical classifiers. They obtained a classification rate of 60%, which was better than that achieved with statistical classifiers.

Dhawan, Chitre, and Moskowitz (1993) used a back-propagation neural network by inputting a set of 10 spatial grey-level dependence features extracted from 85 difficult-to-diagnose mammograms, and obtained a classification accuracy of 74%.

Jiang, Nishikawa, Wolverton, Metz, Schmidt, and Doi (1997) compared the classification performance of a neural network classifier with those of five radiologists. Experimental results showed that the neural network classifier performed better than the radiologists in terms of both the area under the ROC curve (A_z) and the partial area index ($0.90A_z$).

Verma (1998) employed BP with momentum and direct solution method (DSM) — based training algorithms to train a feed-forward neural network for the classification of

micro-calcification. A classification rate of 81.25% was achieved for malignant micro-calcifications.

Christoyianni, Dermatas, and Kokkinakis (1999) compared RBF with MLP networks in the classification of all kind of abnormalities, by processing two types of texture features: a statistical descriptor based on high-order statistics (grey-level histogram moments or GLHM), and the spatial grey-level dependence (SGLD) matrix. Various neural network topologies were tested on the MIAS database. Due to the extensive training and computational complexity for both training and testing, the MLP classifier outperformed the RBF classifier. More specifically, the recognition accuracy of the MLP classifier was approximately 4% better than that obtained by the RBF networks for GLHM features.

Patrocínio, Schiabel, Benatti, Goes, and Nunes (2000) demonstrated that only a few features, such as irregularity, number of micro-calcifications in a cluster, and cluster area, were needed as the inputs of a neural network to separate images into two distinct classes: “suspicious” and “probably benign”. An optimal neural network architecture selected by a simulated annealing optimization technique led to improved classification performance (Gurcan, Chan, Sahiner, Hadjiiski, Petrick, & Helvie, 2002; Gurcan, Sahiner, Chan, Hadjiiski, & Petrick, 2001).

Verma and Zakos (2001) developed a computer-aided diagnosis system for digital mammograms based on fuzzy-neural and feature extraction techniques. They used a fuzzy technique to detect micro-calcification patterns and a BP neural network to classify them. The micro-calcification areas from 40 cases from the Nijmegen digital mammographic database were used for their experiments. A classification rate of over 88.9% was obtained.

Markey, Lo, and Floyd (2002) used a feed-forward back-propagation neural network in their comparative study on the classification of mass and calcification types of breast abnormalities. They attained higher A_z values for mass compared with calcification breast abnormalities on both databases (Duke University Medical Center and The University of Pennsylvania Medical Center). They suggested that different types of abnormalities should be considered separately when evaluating classifier performance.

Gavrielides, Lo, and Floyd (2002) used two feed-forward neural networks, the first to detect if the passed fuzzy-scaled inputs calculated from histogram features of suspicious clusters are positive or negative, and the second to classify the already-detected cluster into “true positive” or “false positive” clusters. They attained a 93.2% classification rate on 44 malignant clusters.

Neural networks in conjunction with evolutionary algorithms have also been used for feature selection and classification of malignant and benign patterns (Anastasio, Yoshida, Nagel, Nishikawa, & Doi, 1998; Lo, Land, & Morrison, 2000; Sahiner, Chan, Petrick, Helvie, & Goodsitt, 1998; Verma & Zakos, 2001; Zhang, Verma, & Kumar, 2004). Lo, Land, and Morrison (1999) proposed an evolutionary algorithm-based neural network for reliable classification of breast lesions into “benign” and “malignant”. The main idea was to reduce the complexity involved in deciding network configuration and learning algorithm in network training, to achieve reliable classification without under/over training of data or network entrapment in local minima — the major risks inherent in the traditional ANN paradigm.

EXPERIMENTAL RESULTS USING NEURAL-EVOLUTIONARY TECHNIQUES

We have recently investigated two novel neural network-based approaches for feature extraction, selection, and classification of breast abnormalities (benign/malignant). The first approach (Zhang, Verma, & Kumar, 2004) is based on neural evolutionary feature selection and classification. The second approach (Panchal & Verma, 2005) is based on a combination of auto-associator and MLP classifier.

A benchmark database (DDSM) was used for the experiments. A total of 126 calcification cases and 117 mass cases (masses may and often do have calcifications in them) were used for the experiments. Experiments were run using 84 calcification cases (43 benign, 41 malignant) for training, and 42 calcification areas (22 benign, 20 malignant) for testing purposes. Another set of experiments was conducted using 78 (40 benign and 38 malignant) mass cases for training, and 39 (20 benign, 19 malignant) for testing.

The following 20 features were used for the experiments: (1) number of pixels, (2) histogram, (3) average grey level, (4) average boundary grey level, (5) difference, (6) contrast, (7) energy, (8) modified energy, (9) entropy, (10) modified entropy, (11) standard deviation, (12) modified standard deviation, (13) skew, (14) modified skew, (15) age, (16) density, (17) calcification type/mass shape, (18) calcification distribution/mass margin, (19) assessment, and (20) subtlety.

Neural Evolutionary Feature Selection and Classification

Evolutionary algorithms for feature selection have been used to find the most significant feature (or combination of features). Various experiments using different parameters were conducted to find the most significant feature (or combination of features) that best classifies a suspicious area as “benign” or “malignant”.

The experiments were conducted by using the classification rate on the test set to calculate the fitness for reproduction in evolutionary feature selection. The number of hidden units and the output threshold were adjusted to find the combination of the features and neural network architecture that achieved the best classification rate on the test set. The results of the experiments are presented below.

Experimental Results for Calcification Cases

The results using 14 features (1-14 listed above) for calcification cases are presented in Table 1. The highest classification rate on the test set was 90.5%. The highest classification rate was obtained with selection of the following feature vector: 10011011000011, which means that the most significant features are numbers 1, 4, 5, 7, 8, 13, and 14. A combination of only seven features provided the best classification rate.

The results using all 20 features for calcification cases are presented in Table 2. The highest classification rate on the test set was 88.1%. The highest classification rate was obtained by selecting the feature vectors at row numbers 5, 7, 8, and 9, which means that there are at least four combinations of features that are more important than others. Feature number one was selected for all four combinations.

Table 1. Feature selection and classification rate using 14 features

Features (1 means selected)	#Hidden units	Classification rate (%)	
		Training set	Test set
10110101000110	4	69.0	81.0
01101000011000	6	73.8	85.7
10100010011111	8	71.4	88.1
10100010011111	8	72.6	85.7
10011011000011	10	76.2	90.5
10011011000011	10	78.6	85.7
11111011011010	12	71.4	85.7
11111011011010	12	78.6	83.3
01100100111001	14	76.2	88.1
10011011000110	14	78.6	83.3
00110110010101	16	73.8	88.1

Table 2. Feature selection and classification rate using 20 features

Features (1 means selected)	#Hidden units	Classification rate (%)	
		Training set	Test set
10001001010010000001	4	73.8	83.3
01100000011110000101	6	75.0	85.7
10010001101011101010	8	76.0	83.3
11110001101011101110	8	72.6	85.7
11110110010000100111	10	63.1	88.1
11110110010000100111	10	81.0	81.0
11111010000111101010	12	73.8	88.1
10101000111011000001	14	79.8	83.3
10101000111011000001	14	72.6	88.1
11010001001100001100	16	69.5	88.1
10001001010010000001	4	73.8	83.3

Experimental Results for Mass Cases

The results using 14 features for mass cases are presented in Table 3. The highest classification rate on the test set was 89.7%. The highest classification rate was obtained with selection of the following feature vector 10100100011010, which means that the most significant features are feature numbers 1, 3, 6, 10, 11, 13. A combination of only six features provided the best classification rate.

The results using 20 features for mass cases are presented in Table 4. The highest classification rate on the test set was 87.2%. The highest classification rate was obtained by selecting the following feature vector: 11000101010011101010, which means that the most significant features are numbers 1, 2, 6, 8, 10, 13, 14, 15, 17, and 19. A combination of only 10 features provided the best classification rate.

Table 3. Feature selection and classification rate using 14 features

Features (1 means selected)	#Hidden units	Classification rate (%)	
		Training set	Test set
10011011000110	4	64.1	84.6
10011011000110	4	70.5	82.1
11100011111010	6	65.4	82.1
11100011111010	6	61.5	84.6
11000110100010	8	70.5	82.1
11000110100010	8	65.4	87.2
11011110000010	10	73.1	82.1
11011110000010	10	65.4	87.2
10100100011010	12	59.0	89.7
10100100011010	12	71.8	82.1
10011011000110	4	64.1	84.6

Table 4. Feature selection and classification rate using 20 features

Features (1 means selected)	#Hidden units	Classification rate (%)	
		Training set	Test set
11110011001010101101	4	60.3	82.1
01110100000110001100	6	60.3	84.6
11100101100001011010	8	70.5	84.6
11100101100001011010	8	75.6	79.4
11001010111010101100	10	66.7	84.6
11001010111010101100	10	74.4	79.5
11000101010011101010	12	75.6	87.2
11000101010011101010	12	78.2	82.1
11110011001010101101	4	60.3	82.1
01110100000110001100	6	60.3	84.6
11100101100001011010	8	70.5	84.6

Auto-Associator Feature Extractor and Neural Classifier

An auto-associator in conjunction with an MLP-based classifier was used to extract new features and classify breast abnormalities into “benign” and “malignant”.

Experimental Results for Calcification Cases

Tables 5 and 6 show the results obtained for calcification cases using 14 and 20 features, respectively. A 100% classification rate on the training set was achieved with both 14 and 20 features. The highest classification rate on the test set reached 90.5% (20 features). With 14 features, the classification rates on the training and test sets were 95.2% and 85.7%, respectively.

Table 5. Classification results with 14 features for calcification cases

Auto-associator		Neural network classifier			
#Hidden units	# Iterations	#Hidden units	# Iterations	Classification rate [%] on training set	Classification rate [%] on test set
4	10000	14	20000	92.9	83.3
6	10000	20	30000	98.8	81
6	20000	14	40000	98.8	83.3
10	10000	18	20000	100	81
12	20000	6	5000	90.5	85.7
14	10000	16	5000	90.5	85.7
14	20000	6	10000	95.2	85.7

Table 6. Classification results with 20 features for calcification cases

Auto-associator		Neural network classifier			
#Hidden units	# Iterations	#Hidden units	# Iterations	Classification rate [%] on training set	Classification rate [%] on test set
8	10000	6	20000	100	81
8	10000	8	20000	100	83.3
8	10000	8	40000	100	85.7
10	10000	6	5000	100	83.3
10	20000	14	30000	100	88.1
10	30000	10	5000	100	88.1
10	70000	6	10000	100	90.5

Table 7. Classification results with 14 features for mass cases

Auto-associator			Neural network classifier		
#Hidden units	# Iterations	#Hidden units	# Iterations	Classification rate [%] on training set	Classification rate [%] on test set
4	10000	6	5000	78.2	82.1
4	10000	10	5000	75.6	84.6
4	10000	10	20000	80.8	87.2
4	10000	14	30000	78.2	89.7
6	20000	8	50000	100	74.4
6	40000	12	20000	84.6	84.6
6	70000	20	20000	83.3	89.7
6	100000	6	50000	87.2	82.1

Table 8. Classification results with 20 features for mass cases

Auto-associator				Neural network classifier	
#Hidden units	# Iterations	#Hidden units	# Iterations	Classification rate [%] on training set	Classification rate [%] on test set
4	10000	10	30000	100	66.7
4	20000	8	20000	87.2	69.2
4	30000	8	20000	85.9	71.8
4	30000	10	50000	97.4	71.8
10	20000	6	30000	100	69.2
12	30000	8	20000	100	76.9

Experimental results for mass cases

The results for mass cases are presented in Tables 7 and 8. The highest classification rate with 14 features on the test set was 89.7%. The classifier network was trained using 20 hidden neurons and 20,000 iterations on neural associative patterns, obtained from an auto-associator neural network (AANN) trained with six hidden neurons over 70,000 iterations.

The highest classification rate with 20 features on the test set was 76.9%. The classifier network was trained using eight hidden neurons and 20,000 iterations on neural associative patterns, obtained from an AANN trained with 12 hidden neurons over 30,000 iterations.


Table 9 shows the distribution of both correct “malignant” and “benign” classification as well as misclassification results (in other words, “malignant” classified as

Table 9. Distribution of malignant and benign classification and misclassification results

Dataset	Feature* vector	Classification rate (%)				
		Total test	Malignant (TPF)	Malignant as benign (FNF)	Benign (TNF)	Benign as malignant (FPF)
<i>Figure 4 – Y axis scale</i>		<i>(1)</i>	<i>(2)</i>	<i>(3)</i>	<i>(4)</i>	<i>(5)</i>
Calcification	1	85.7	76.2	23.8	95.2	4.8
Calcification	2	90.5	85.7	14.3	95.2	4.8
Mass	1	89.7	95.2	4.8	83.3	16.7
Mass	2	76.9	66.7	33.3	88.9	11.1

*Feature vector 1-14 features; Feature vector 2-20 features

Figure 4. Classification results with feature vector 1 and 2



“benign” and vice versa), for the highest test classification rates obtained using the proposed research technique. Major changes were observed in malignant classification for both groups of the dataset when both types of features were combined together. For calcification cases, the malignant classification rate improved when 14 features (1-14, as listed at the start of this section) were combined with six other features (15-20 listed earlier); the benign classification rate remained unchanged. For mass cases, when both types of features were combined, the malignant classification rate decreased significantly; in contrast, benign classification improved slightly.

Table 10. Comparative analysis of neural network-based techniques in digital mammography

Neural technique	Feature extraction	Feature selection	Databases	Classification rate [%]	Reference
SOM-MLP	SOM	N/A	Duke University DDSM	98 (sensitivity) 25 (specificity) 86.5	Markey et al. (2003)
RBFNN	Texture	N/A	MIAS	65 and 77	Bovis et al. (2000)
BPNN	Cluster	N/A	unknown	60	Chitre et al. (1993)
BPNN	Wavelets	Neural network	unknown	80	Kocur et al. (1996)
RBFNN	Texture & wavelets	N/A	unknown Database	74	Dhawan et al. (1995)
BPNN	N/A	N/A	157 images	90	Jiang et al. (1997)
BPNN	ROI's	N/A	LLN	87	Verma (1998)
DSMNN	ROI's	N/A	Nijmegen	86.6	Verma (1998)
BPNN	14 features	Neural network	Nijmegen	88.9	Verma & Zakos (2001)
EANN	14 features	EA	DDSM	90.5	Zhang et al. (2004)
AANN	14 features and AANN	N/A	DDSM	90.5	Panchal & Verma (2005)
BPNN	BI-RADS	N/A	Duke University	Mass Az – 0.93 Calcification Az – 0.63	Markey, Lo, & Floyd (2002)
			University of Pennsylvania	Mass Az – 0.88 Calcification Az – 0.76	
BPNN	4 features	N/A	MIAS	85	Lee & Tsai (2004)
GANN	4 features	N/A	MIAS	77	Lee & Tsai (2004)
BPNN	LVQ-NN	N/A	DDSM	CC View – 97.3 MLO View - 96.18	Khuwaja & Abu-Rezq (2004)
EPPNN	BI-RADS	EPNN	Duke University	74	Lo et al. (2000)
PNN	BI-RADS	N/A	Duke University	73.5, Az Value – 0.826	Lo et al. (2000)
BPNN	Shape	N/A	Foothills Hospital	90.5 (94-Benign, 87-Malignant)	Shen et al. (1993)
RBFNN	Texture	N/A	MIAS	78.15	Christoyianni et al. (2000)
MLPNN	Texture	N/A	MIAS	82.35	Christoyianni et al. (1999)
BPNN	Histogram	Fuzzy	DDSM	93.2 - Malignant	Gavrielides et al. (2002)
BPNN	Texture	N/A	MIAS	77	Bovis & Singh (2000)

COMPARATIVE ANALYSIS

In the last few decades most researchers in the area of digital mammography have used back-propagation neural networks (BPNN) with various feature extraction and selection techniques. However, a few researchers have applied other types of neural networks to the classification of breast abnormalities in digital mammograms, including: radial basis function (RBFNN), learning vector quantization (LVQNN), probabilistic (PNN), evolutionary programming and probabilistic (EPPNN), Kohonen networks/ self-organising maps (SOM), direct solution method-based (DSMNN), evolutionary algorithm-based (EANN) and auto-associator-based (AANN).

Results obtained using these nine different neural techniques are presented in Table 10 (taken from published papers). It is easy to compare and analyse the results if they are produced using the same criteria and benchmark database; however, this is not the case in using neural networks in digital mammography. Some researchers combine detection and classification, while others use very small databases. Also, some researchers are not very clear about their results in terms of classification rates for benign and malignant patterns. Thus there is no consistency in conducting experiments and reporting results, which makes it very difficult to compare them and draw conclusions as to which technique is best.

Despite the popularity of benchmark databases such as DDSM and MIAS, not all researchers list results obtained from the same test set. The DDSM database, for instance, has freely available training and test sets. Recently some researchers using this database report results appropriately by listing the classification rates on the test set; they list the classification rates for benign and malignant cases, and some also included true positive, false positive, and ROC curves.

It is surprising that RBF neural networks do not perform well, achieving only 78.15% classification using texture features on the MIAS database. SOMs in conjunction with MLPs perform better, achieving 98% sensitivity and 25% specificity (which corresponds to a classification rate of 86.5%). Back-propagation neural networks do not achieve consistent results, the lowest (60%) being achieved with clustering, and the highest (93.7%) with a histogram feature and fuzzy selection. Many other results lay in between, a typical one being 90.5% using auto-associators or evolutionary algorithms. We have some concerns about the high results (over 97%) reported using BPNN and LVQ-NN. Overall, BPNN in conjunction with various feature extraction and selection techniques performs well.

DISCUSSION AND CONCLUSIONS

We have presented an application of artificial neural networks for the classification of breast abnormalities in digital mammograms. A review of ANN-based techniques for detection, feature extraction, feature selection, and classification of “benign” and “malignant” patterns has been presented. Some recent advances in feature extraction and selection in conjunction with neural classifiers have been described, and experimental results using the DDSM database were presented in Tables 1-9. The highest classification rate achieved was 90.5%.

It was found that combining features leads to higher classification rates than for separate features. Auto-associator-based feature extraction (without using a separate selection strategy) performed as well as neural evolutionary selection, both achieving a 90.5% classification rate for calcification data.

A comparative analysis of neural network-based techniques using different databases, feature extraction, feature selection, and classification techniques has been presented. Nine different types of neural networks have been applied to the classification of benign and malignant patterns in digital mammograms. The most frequently used neural networks for such classification were based on the back-propagation algorithm. Neural networks for feature extraction and selection have also been based on self-organization, auto-association, and evolutionary algorithms. The classification accuracy achieved by neural network-based techniques varies from 60% to 97.3%. The reasons for such variation were the various feature extraction and selection techniques themselves, the training parameters, and the database used. Many researchers have used very small databases and different criteria to report their final results, making it very difficult to compare and analyse the classification rates obtained using these different neural techniques. Therefore, we would like to encourage researchers to report their results on a *benchmark* database test set, in terms of classification rates for benign and malignant micro-calcifications, as well as listing all training parameters.

Finally, we would like to say that the review of advances in neural techniques and the results presented in this chapter show that neural network-based techniques are a promising tool for the classification of benign and malignant patterns in digital mammograms.

REFERENCES

- Aghdasi, F., Ward, R., & Palcic, B. (1994). Restoration on mammographic images in the presence of signal-dependent noise. In K. W. Bowyer & S. Astley (Eds.), *State of the art in digital mammographic image analysis* (Vol. 7(6), pp. 42-63). Singapore: World Scientific Publishing.
- Anastasio, M., Yoshida, H., Nagel, R., Nishikawa, R., & Doi, K. (1998). A genetic algorithm-based method for optimizing the performance of a computer-aided diagnosis scheme for detection of clustered micro-calcifications in mammograms. *Medical Physics*, 25(9), 1613-1620.
- Bakic, P., & Barzakovic, D. (1997, September 8-9). Application of neural networks in computer aided diagnosis of breast cancer. *Proceedings of the 5th Workshop on Neural Network Applications — NEUREL*, Belgrade, Yugoslavia.
- Blake, L., & Merz, J. (1998). *UCI Repository of Machine Learning Databases*. Retrieved from the University of California Web site at <http://www.ics.uci.edu/~mllearn/MLRepository.html>
- Bovis, K., & Singh, S. (2000, September 3-8). Detection of masses in mammograms using texture measures. *Proceedings of the 15th IEEE International Conference on Pattern Recognition*, Quebec, Canada (Vol. 2, pp. 267-270).
- Bovis, K., Singh, S., Fieldsend, J., & Pinder, C. (2000, July 24-27). Identification of masses in digital mammograms with MLP and RBF nets. *Proceedings of the IEEE International Joint Conference on Neural Networks*, Como, Italy (Vol. 1, pp. 1342-1348).

- Chan, H., Sahiner, B., & Lam, K. (1998). Computerized analysis of mammographic micro-calcifications in morphological and texture features space. *Medical Physics*, 25(10), 2007-2019.
- Cheng, H., Cai, X., Chen, X., Hu, L., & Lou, X. (2003). Computer-aided detection and classification of micro-calcifications in mammograms: A survey. *Pattern Recognition*, 36, 2967-2991.
- Chitre Y., Dhawan, A. P., & Moskowitz, M. (1993). Artificial neural network based classification of mammographic micro-calcifications using image structure features. *International Journal of Pattern Recognition and Artificial Intelligence*, 7, 1377-1401.
- Christoyianni, L., Dermatas, E., & Kokkinakis, G. (1999, September 5-8). Neural classification of abnormal tissue in digital mammography using statistical features of the texture. *Proceedings of the IEEE International Conference on Electronics Circuits and Systems*, Pafos, Cyprus (Vol. 1, pp. 117-120).
- Christoyianni, L., Dermatas, E., & Kokkinakis, G. (2000). Fast detection of masses in computer-aided mammography. *IEEE Signal Processing Magazine*, 17, 54-64.
- Dhawan, A., Chitre, Y., Bonasso, C., & Wheeler, K. (1995, September 20-23). Radial basis function based classification of mammographic micro-calcifications using texture features. *Proceedings of the IEEE Engineering in Medicine and Biology 17th Annual Conference & 21st Canadian Medical and Biological Engineering Conference*, Montreal, Canada (pp. 535-536).
- Dhawan, A., Chitre, Y., & Moskowitz, M. (1993). Artificial neural network based classification of mammographic micro-calcifications using image structure features. In R. S. Acharya & D. B. Goldgof (Eds.), *Biomedical Image Processing and Biomedical Visualization — Proceedings of SPIE (the International Society of Optical Engineering)* (Vol. 1905, pp. 820-831).
- Diahi, J., Frouge, C., Giron, A., & Fertil, B. (1996, June 9-12). Artificial neural networks for detection of breast cancer in mammography. *Proceedings of the 3rd International Workshop on Digital Mammography*, Chicago (pp. 329-334).
- Erich, S., & Ferdinand, S. (2000). Computer aided diagnosis of clustered micro-calcifications using artificial neural nets. *Journal of Computing and Information Technology*, 8, 151-160.
- Gavrielides, M. A., Lo, J. Y. Jr., & Floyd, C. E. (2002). Parameter optimization of a computer-aided diagnosis scheme for the segmentation of micro-calcification clusters in mammograms. *Medical Physics*, 29, 475-483.
- Gonzalez, R., & Woods, R. (1993). *Digital image processing*. Reading, MA: Addison Wesley.
- Gurcan, M., Chan, H., Sahiner, B., Hadjiiski, L., Petrick, N., & Helvie, M. (2002). Optimal neural network architecture selection: Improvement in computerized detection of micro-calcifications. *Academic Radiology*, 9, 420-429.
- Gurcan, M., Sahiner, B., Chan, H., Hadjiiski, L., & Petrick, N. (2001). Selection of an optimal neural network architecture for computer-aided detection of micro-calcifications - Comparison of automated optimization techniques. *Medical Physics*, 28, 1937-1948.
- Heath, M., Bowyer, K., Kopans, D., Moore, R., & Kegelmeyer, P. (2000, June 11-14). The digital database for screening mammography. In M. J. Yaffe (Ed.), *Proceedings of*

- the 5th International Workshop on Digital Mammography*, Toronto, Canada. Madison, WI: Medical Physics Publishing.
- Jain, A. (1995). *Fundamentals of digital image processing*. Upper Saddle River, NJ: Prentice-Hall.
- Jiang, Y., Nishikawa, R., Wolverton, D., Metz, C., Schmidt, R., & Doi, K. (1997, October 30-November 2). Computerized classification of malignant and benign clustered micro-calcifications in mammograms. *Proceedings of the 19th International Conference of the IEEE/Engineering in Medicine and Biology Society* (pp. 521-523).
- Karssemeijer, N. (1994). Adaptive noise equalization and recognition of micro-calcification clusters in mammograms. In S. Astley (Eds.), *State of the art of digital mammographic image analysis* (Vol. 9(8), pp. 148-166). Singapore: World Scientific Publishing Series in Machine Perception and Artificial Intelligence.
- Karssemeijer, M., Thijssen, M., Hendriks, J., & van Erning, L. (Eds.). (1998). *Digital mammography* (pp. 237-244). Nijmegen: Kluwer Academic Publishers.
- Khuwaja, G. A., & Abu-Rezq, A. N. (2004). Bi-model breast cancer classification system. *Pattern Analysis Applications*, 7, 235-242.
- Kim, J. K., & Park, H. W. (1999). Statistical textural features for detection of micro-calcifications in digitized mammograms. *IEEE Transactions on Medical Imaging*, 18, 231-238.
- Kocur, C., Gogers, S., Myers, L., Burns, T., Kabrisky, M., HoImeister, J., Baver, K., & Steppe, J. (1996). Using neural networks to select wavelet features for breast cancer diagnosis. *IEEE Engineering in Medicine and Biology Magazine*, 15(3), 95-102.
- Kopans, D. (1998). *Breast imaging* (2nd ed.). Philadelphia: Lippincott-Raven.
- Lee, G., & Bottema, M. (2001, June). Classification of masses in screening mammograms as benign or malignant. In M. J. Yaffe (Ed.), *Proceedings of the 5th International Workshop on Digital Mammography — IWDM2000*, Toronto, Canada (p. 52). Madison, WI: Medical Physics Publishing.
- Lee, Y., & Tsai, D. (2004). Computerized classification of micro-calcifications on mammograms using fuzzy logic and genetic algorithm. In M. Sonka & J. M. Fitzpatrick (Eds.), *Medical Imaging — Proceedings of SPIE* (Vol. 5370, pp. 952-959).
- Lo, J. Y., Gavrielides M., Markey M. K., & Esneck J. L. (2003). Computer-aided classification of breast micro-calcification clusters: Merging of features from image processing and radiologists. In M. Sonka & J. M. Fitzpatrick (Eds.), *Medical Imaging — Proceedings of SPIE* (Vol. 5032, pp. 882-889).
- Lo, J. Y., Land, W. H., & Morrison, C. T. (1999, July 10-16). Application of evolutionary programming and probabilistic neural networks to breast cancer diagnosis. *Proceedings of the IEEE International Joint Conference on Neural Networks*, Washington, DC (Vol. 5, pp. 3712-3716).
- Lo, J. Y., Land, W. H., & Morrison, C. T. (2000). Evolutionary programming technique for reducing complexity of artificial neural networks for breast cancer diagnosis. In K. M. Hanson (Ed.), *Medical Imaging, Image Processing — Proceedings of SPIE* (Vol. 3979, pp. 153-158).
- Markey, M. K., Lo, J. Y., & Floyd, C. E. (2002). Differences between computer-aided diagnosis of breast masses and that of calcifications. *Radiology*, 223, 489-493.
- Markey, M., Lo, J., Tourassi, G., & Floyd, C. (2003). Self-organizing map for cluster analysis of a breast cancer database. *Artificial Intelligence in Medicine*, 27, 113-127.

- Masek, M., Attikiouzel, Y., & deSilva, C. J. S. (2001). Automatic removal of high intensity labels and noise from mammograms. In M. J. Yaffe (Ed.), *Proceedings of the 5th International Workshop on Digital Mammography — IWDM2000*, Toronto, Canada, June 2000 (pp. 576-581). Madison, WI: Medical Physics Publishing.
- Neiber, H., Mueller, T., & Stotzka, R. (2001). Local contrast enhancement for the detection of micro-calcifications. In M. J. Yaffe (Ed.), *Proceedings of the 5th International Workshop on Digital Mammography — IWDM2000*, Toronto, Canada, June 2000, (p. 103). Madison, WI: Medical Physics Publishing.
- Panchal, R., & Verma, B. (2005, August 1-4). Classification of breast abnormalities in digital mammograms using image and BI-RADS features in conjunction with neural network. In *Proceedings of the IEEE International Joint Conference on Neural Networks*, Montreal, Canada.
- Papadopoulos A., Fotiadis D. I., & Likas, A. (2002). An automatic micro-calcification detection system based on a hybrid neural network classifier. *Artificial Intelligence in Medicine*, 25, 149-167.
- Patrocínio, A., Schiabel, H., Benatti, R., Goes, C., & Nunes, F. (2000, July 23-28). Investigation of clustered micro-calcification features for an automated classifier as part of a mammography CAD scheme. *Proceedings of the 22nd Annual IEEE/Engineering in Medicine and Biology Society International Conference*, Chicago (pp. 1203-1205).
- Pereira, R., & Azevedo, M. (2001). Wavelet transform improves contrast enhancement of digital mammographic images. In M. J. Yaffe (Ed.), *Proceedings of the 5th International Workshop on Digital Mammography — IWDM2000*, Toronto, Canada, June 2000 (p. 104). Madison, WI: Medical Physics Publishing.
- Petrick, N., Chan, H., Sahiner, B., Wei, D., Helvie, M., Goodsitt, M., & Adler, D. (1995, July 17-21). Automated detection of breast masses on digital mammograms using a convolution neural network for morphological and texture classification. *Proceedings of the World Congress on Neural Networks*, Washington, DC (Vol. 2, pp. 872-875).
- Sahiner, B., Chan, H., Petrick, N., Helvie, M., & Goodsitt, M. (1998). Design of a high-sensitivity classifier based on a genetic algorithm: Application to computer-aided diagnosis. *Physics in Medicine and Biology*, 43, 2853-2871.
- Sajda, P., Spence, C., & Pearson, J. (1995, October 23-26). A hierarchical neural network architecture that learns target context: Applications to digital mammography. *Proceedings of the 2nd IEEE International Conference on Image Processing*, Washington, DC (pp. 3149-3151).
- Sajda, P., Spence, C., Pearson, J., & Nishikawa, R. (1996, June 9-12). Integrating multi-resolution and contextual information for improved micro-calcification detection in CAD. In K. Doi, M. L. Giger, R. M. Nishikawa, & R. A. Schmidt (Eds.), *Proceedings of the 3rd IEEE International Workshop on Digital Mammography*, Chicago (pp. 287-290). Amsterdam, The Netherlands: Elsevier.
- Shen, L., Rangayyan, R. M., & Desautels, J. E. L. (1993). Detection and classification of mammographic calcifications. *International Journal of Pattern Recognition and Artificial Intelligence*, 7, 1403-1416.
- Sonka, M., & Fitzpatrick, J. (2000). *Medical image processing and analysis*. Bellingham, WA: SPIE Press.

- Umbaugh, S. (1998). *Computer vision and image processing: A practical approach using CVPtools*. Upper Saddle River, NJ: Prentice Hall.
- Verma, B. (1998, May 4-9). A neural network based technique to locate and classify micro-calcifications in digital mammograms. *Proceedings of the IEEE International Joint Conference on Neural Networks*, Anchorage, Alaska (Vol. 3, pp. 1790-1793).
- Verma, B., & Zakos, J. (2001). A computer-aided diagnosis system for digital mammograms based on fuzzy-neural and feature extraction techniques. *IEEE Transactions on Information Technology in Biomedicine*, 5, 46-54.
- Wei, Q., Laurence, P., & Clark, L. (1994, November). Adaptive multistage nonlinear filtering and wavelet for medical image enhancement. *Proceedings of the IEEE International Conference on Image Processing — ICIP*, 13-16, Austin, TX (Vol. 3, pp. 711-715).
- Wei, D., Nishikawa, R., & Doi, K. (1996, June 9-12). On the application of a shift invariant feed-forward artificial neural network for detection of clustered micro-calcifications. In K. Doi, M. L. Giger, R. M. Nishikawa, & R. A. Schmidt (Eds.), *Proceedings of the 3rd IEEE International Workshop on Digital Mammography*, Chicago (pp. 283-286). Amsterdam, The Netherlands: Elsevier.
- Woods, K., Doss, C., Bowyer, K., Solka, J., Priebe, C., & Kegelmeyer, W. (1993). Comparative evaluation of pattern recognition techniques for detection of micro-calcifications in mammography. *International Journal of Pattern Recognition and Artificial Intelligence*, 7(6), 1417-1436.
- Wroblewska, A., Boninski, P., Przelaskowski, A., & Kazubek, M., (2003). Segmentation and feature extraction for reliable classification of micro-calcification in digital mammograms. *Opto-Electronics Review*, 11, 227-235.
- Wu, Y. (1993). Application of neural networks in mammography. *Radiology*, 187, 81-87.
- Wu, Y., He, J., Man, Y., & Arribas, J. I. (2004, July 25-29). Neural network fusion strategies for identifying breast masses. In *Proceedings of the IEEE International Joint Conference on Neural Networks*, Budapest, Hungary (Vol. 3, pp. 2437-2442).
- Yin, F., Giger, M., Vyborny, C., Doi, K., & Schmidt, R. (1993). Comparison of bilateral-subtraction and single-image processing techniques in the computerised detection of mammographic masses. *Investigative Radiology*, 28(6), 473-481.
- Yoon, J., Ro, Y., Kim, S., & Park, D. (2002). Contrast enhancement of mammography image using homomorphic filter in wavelet domain. *IEICE Transactions on Information & Systems*, E85-D(1), 298-303.
- Yu, S., & Guan, L. (2000). A CAD system for the automatic detection of clustered micro-calcifications in digitized mammogram films. *IEEE Transactions on Medical Imaging*, 19, 115-126.
- Zhang, P., Verma, B., & Kumar, K. (2004, July 25-29). A neural-genetic algorithm for feature selection and breast abnormality classification in digital mammography. *Proceedings of the IEEE International Joint Conference on Neural Networks*, Budapest, Hungary (Vol. 3, pp. 2303-2309).
- Zheng, B., Qian, W., & Clarke, L. (1994, June 27-July 2). Multistage neural network for pattern recognition in mammogram screening. *Proceedings of the IEEE International Joint Conference on Neural Networks*, Orlando, FL (pp. 3437-3447).

Chapter X

Swarm Intelligence and the Taguchi Method for Identification of Fuzzy Models

Arun Khosla, National Institute of Technology, Jalandhar, India

Shakti Kumar, Haryana Engineering College, Jagadhari, India

K. K. Aggarwal, GGS Indraprastha University, Delhi, India

ABSTRACT

Nature is a wonderful source of inspiration for building models and techniques for solving difficult problems in design, optimisation, and control. More specifically, the study of evolution, the human immune system, and the collective behaviour of insects/birds have guided the origin of evolutionary algorithms, artificial immune systems, and optimisation techniques based on swarm intelligence, respectively. In this chapter, we present the use of particle swarm optimisation (PSO) and the Taguchi method for the identification of optimised fuzzy models from the available data. PSO is a member of the broad category of swarm intelligence (SI) techniques based on the metaphor of social interaction. It has been used for finding promising solutions in complex search spaces through the interaction of particles in a swarm, and is especially useful when dealing

with a high number of dimensions and situations where problem-specific information is not available. However, caution needs to be exercised in selecting PSO, as the performance of PSO largely depends on their values. In this chapter, a systematic reasoning approach based on the Taguchi method is also presented to quickly identify PSO parameters. The Taguchi method is a robust design approach that helps in optimisation, and which requires relatively few experiments. Although we focus here on the use of PSO and the Taguchi method for fuzzy model identification, these techniques have much broader use and application. In order to validate our approach, data from the rapid Nickel-Cadmium (Ni-Cd) battery charger developed by the authors were used. The results are based on real data and illustrate the viability and efficiency of the approach.

INTRODUCTION

In recent years, the concept of fuzzy set theory has received considerable attention, both in academia and industry, due to its ability to handle ambiguous or vague concepts of human perception for complex systems problems, in which it is extremely difficult to describe the system models mathematically. Moreover, the concept of fuzzy set has been applied successfully in many disciplines.

The problem of fuzzy system modelling or fuzzy model identification is generally the determination of a fuzzy model for a system or process by making use of linguistic information obtained from human experts, and/or numerical information obtained from input-output measurements. The former approach is known as knowledge-driven modelling, while the latter is known as data-driven modelling. In this chapter, attention is focused on building fuzzy models from the available data using PSO, a relatively new optimisation technique.

The performance of PSO and other evolutionary algorithms, to a great extent, depends upon the choice of appropriate parameters. Generally, these parameters are selected through a hit-and-miss (trial-and-error) process, which is unsystematic and requires unnecessarily rigorous experimentation. In this chapter, we propose a systematic approach based on the Taguchi method for the identification of the optimal strategy parameters of PSO for fuzzy model identification.

The remainder of the chapter is structured as follows: The next section serves as a brief introduction to swarm intelligence and the PSO algorithm. Brief information about the fuzzy model identification problem is provided in the following section. Some details about the rapid Ni-Cd battery charger are then provided. A framework for fuzzy model identification through using the PSO algorithm is presented in the next section. The proposed framework has been applied to identify fuzzy models for a rapid Ni-Cd battery charger. A description of the Taguchi method, together with selection of appropriate parameters for the PSO algorithm for fuzzy model identification, is given in the next-to-last section. Here we also present a comparison of the computational efforts required by both the Taguchi method and the traditional approach, which involves exhaustive combinations of the PSO operating parameters. Concluding remarks are made in the final section.

SWARM INTELLIGENCE AND THE PSO ALGORITHM

The field of swarm intelligence (SI) has been inspired by the social behaviour of ants, termites, bees, wasps, birds, fishes, and other biological creatures, and is emerging as an innovative and powerful computational metaphor for solving complex problems in design, optimisation, and control. SI could be defined as (Kennedy & Eberhart, 2001):

Any attempt to design algorithms or distributed problem-solving devices inspired by the collective behaviour of insect colonies and other animal societies.

Research into social insect behaviour suggests that intelligent group behaviour emerges out of simple interactions between individuals, which otherwise have limited capabilities.

The motivation behind the PSO algorithm is the social behaviour of animals, for instance, the flocking of birds and the schooling of fish. PSO has its origin in simulation for visualising the synchronized choreography of bird flocks by incorporating concepts such as nearest-neighbour velocity matching and acceleration by distance (Eberhart & Shi, 2001; Kennedy & Eberhart, 1995, 2001; Parsopoulos & Vrahatis, 2002). Later, it was realized that such simulation could be used for optimisation and resulted in the first simple version of PSO (Kennedy & Eberhart, 1995). Since then, many variants of PSO have been suggested by different researchers (Eberhart & Kennedy, 1995; Shi & Eberhart, 2001; Xie, Zhang, & Yang, 2002).

The PSO algorithm, like other evolutionary algorithms (EAs), is a stochastic technique that uses a population of potential solutions (called particles) to probe the search space and also does not require gradient information of the objective function under consideration. Unlike other EAs, such as genetic algorithm (GA), PSO favours *collaboration* among the candidate solutions instead of *rivalry*. The majority of evolutionary optimisation algorithms are based on Darwin's theory of "Survival-of-the-Fittest". In particle swarms, it is not the drive of survival that breeds quality solutions; rather, the individuals strive to improve themselves by imitating traits from their successful peers. In PSO, the particles have an adaptable velocity that determines their movement in the search space. Each particle also has a memory and hence is capable of remembering the best position in the search space it has ever visited.

One of the most promising advantages of PSO over GAs is its algorithmic simplicity, as it uses only primitive mathematical operators, which accounts for its low computational overhead. The three major operators in GAs are: selection, crossover, and mutation, and there exist many options for implementation thereof. For example, one may go for roulette-wheel or tournament selection, single or double-point crossover, and so on, whereas in PSO the *only* operation is velocity calculation.

Two broad variants of PSO algorithm have been developed: one with a global neighbourhood called the *gbest* model, and the other with a local neighbourhood known as the *lbest* model. The *gbest* model maintains only a single best solution, and each particle moves towards its previous best position and towards the best particle in the whole swarm. The best particle acts as an attractor, pulling all the particles towards it.

In the *lbest* model, each particle moves towards its previous best position, and also towards the best particle in its restricted neighbourhood and thus maintains *multiple* attractors. A sub-set of particles is defined for each particle from which the local best particle is then selected. Particles selected to be in a sub-set have no relation to each other in the search space domain. Thus, the difference between two variants is based on the set of particles with which a given particle will interact directly. Note that the *gbest* model is actually a special case of *lbest* (in which neighbourhood size is equal to swarm size). In this chapter, we use the *gbest* model for fuzzy model identification.

Consider that the search space is d -dimensional and i th particle in the swarm can be represented by $X_i = (x_{i1}, x_{i2}, \dots, x_{id})$ and its velocity can be represented by another d -dimensional vector $V_i = (v_{i1}, v_{i2}, \dots, v_{id})$. Let the best previously-visited position of this particle be denoted by $P_i = (p_{i1}, p_{i2}, \dots, p_{id})$. If the g th particle is the best particle and the iteration number is denoted by the superscript, then the swarm is modified according to the following equations (1) and (2) suggested by Eberhart and Shi (2001).

$$v_{id}^{n+1} = \chi(wv_{id}^n + c_1r_1^n(p_{id}^n - x_{id}^n) + c_2r_2^n(p_{gd}^n - x_{id}^n)) \quad (1)$$

$$x_{id}^{n+1} = x_{id}^n + v_{id}^{n+1} \quad (2)$$

where χ = constriction factor; w = inertia weight; c_1 = cognitive parameter; c_2 = social parameter, and r_1 and r_2 are random numbers distributed evenly in the range (0,1).

These parameters, χ , w , c_1 , and c_2 , along with V_{max} are the strategy/operating parameters, and the performance of the PSO algorithm to a great extent depends on appropriate selection of these parameters (Eberhart & Shi, 2001). The parameter V_{max} is defined by the user to be the maximum velocity along any dimension, which implies that if the velocity along any dimension exceeds V_{max} , it shall be clamped to this value. The inertia weight governs how much of the previous velocity should be retained from the previous time step. Generally the inertia weight is not kept fixed, but is varied as the algorithm progresses, so as to improve performance (Eberhart & Shi, 2001; Parsopoulos & Vrahatis, 2002). This setting allows the PSO to explore a large area at the start of a simulation run, and to refine the search later by a smaller inertia weight. The parameters c_1 and c_2 influence the maximum size of the step that a particle can take in a single iteration, and the random numbers r_1 and r_2 help in maintaining diversity of the population. The constriction factor was introduced by Clerc (1999) to ensure convergence.

THE FUZZY MODEL IDENTIFICATION PROBLEM

Generally, the problem of fuzzy model identification includes the following issues (Hellendoorn & Driankov, 1997; Yen & Langari, 2003):

- Selecting the type of fuzzy model,
- Selecting the input and output variables for the model,
- Identifying the structure of the fuzzy model, which includes determination of the number and types of membership functions for the input and output variables, as well as the number of fuzzy rules,

- Identifying the parameters of antecedent and consequent membership functions, and
- Identifying the consequent parameters of the fuzzy rule base.

Three commonly-used types of fuzzy model are Mamdani-type, Takagi-Sugeno, and Singleton.

In Mamdani models, each fuzzy rule is of the form:

$$R_i: \text{If } x_{i1} \text{ is } A_{i1} \text{ and } \dots \text{ and } x_{in} \text{ is } A_{in} \text{ then } y \text{ is } B \quad (3)$$

In Takagi-Sugeno models, each fuzzy rule is of the form:

$$R_i: \text{If } x_{i1} \text{ is } A_{i1} \text{ and } \dots \text{ and } x_{in} \text{ is } A_{in} \text{ then } y \text{ is } \sum_{i=1}^n a_i x_i + C \quad (4)$$

Whereas in Singleton models, each fuzzy rule is of the form:

$$R_i: \text{If } x_{i1} \text{ is } A_{i1} \text{ and } \dots \text{ and } x_{in} \text{ is } A_{in} \text{ then } y \text{ is } C \quad (5)$$

where x_1, \dots, x_n are the input variables and y is the output variable, A_{i1}, \dots, A_{in}, B are the linguistic values of the input and output variables in the i th fuzzy rule, and a_i and C are constants. In fact, the Singleton fuzzy model can be seen as a special case of the Takagi-Sugeno model when $a_i=0$. The input and output variables take their values in their respective universes of discourse or domains. In this chapter, we have considered identification based on Mamdani and Singleton fuzzy models only.

Some commonly used techniques for creating fuzzy models from the available input-output data are Genetic Algorithms (Bastian, 1996; Carse, Fogarty, & Munro, 1996; Nelles, 1996; Nozaki, Morisawa, & Ishibuchi, 1996), the fuzzy c-means (FCM) clustering algorithm (Khosla, Kumar, & Aggarwal, 2003a; Setnes & Roubos, 1999), neural networks (Hellendoorn & Driankov, 1997), and the adaptive neuro fuzzy inference system model (ANFIS) (Khosla, Kumar, & Aggarwal, 2003b; Melin & Castillo, 2005).

RAPID NICKEL-CADMIUM BATTERY CHARGER

Batteries can be classified into two main groups: primary batteries and secondary batteries. Unlike primary batteries, secondary batteries once discharged can be returned to their fully charged state and can be discharged and charged many times, thus making them economical. Nickel-cadmium (Ni-Cd), nickel metal hydride (Ni-MH) and lithium ion (Li-Ion) are some commonly-used secondary batteries.

The most common method to charge Ni-Cd batteries is by means of a constant-current source at the rate of 0.1C (trickle charge), where the charging rate is commonly expressed as a multiple of the rated capacity of the battery (Linden, 1995); for example, for a battery with $C=500$ mAh, 0.1C corresponds to a charging current of 50 mA. At this rate, the battery takes between 12 and 16 hours to charge, and can withstand overcharg-

Table 1. Input and output variables for rapid Ni-Cd battery charger, along with their universes of discourse

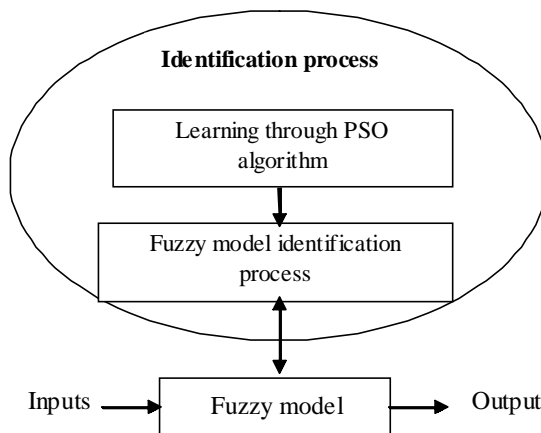
Input variables	Universe of discourse
Temperature (T)	0-50°C
Temperature gradient (dT/dt)	0-1 (°C/sec)
Output variable	
Charging rate (Ct)	0-8C

ing without harm. Some chargers have the capability of charging batteries in about 5 hours using higher charging currents. However, with high charging rates (C/3 or higher), care must be taken to avoid overcharging, as it may result in excessive rise in temperature, which can harm the batteries (Buchmann, 1997). The main objective for the development of a rapid battery charger is to charge Ni-Cd batteries quickly, but without causing any damage to them. Since the behaviour of Ni-Cd batteries at very high charging rates is not available, there was need to obtain these through experimentation. Based on initial trials with a charging rate of 8C, coupled with the fact that a Ni-Cd battery is capable of supplying currents of the order of 8C without damage (Wan, 1996), the upper limit of the charging current was fixed at 8C — namely 4A — since the target batteries had a capacity C of 500 mAh.

Based on rigorous experimentation with Ni-Cd batteries (Khosla, 1997; Khosla, Kumar, & Aggarwal, 2002), it was observed that the two input variables used to control the charging rate (Ct) are absolute temperature (T) of the battery and its temperature gradient (dT/dt). The input and output variables identified for a rapid Ni-Cd battery charger along with their universes of discourse are listed in Table 1 (this data set comprised 561 points, and is available at <http://research.4t.com>).

A FRAMEWORK FOR FUZZY MODEL IDENTIFICATION WITH THE PSO ALGORITHM

Optimisation plays an important role in many fields. A common problem, however, is model fitting, where the goal is to find the model parameters so that the error between the desired output and actual output can be minimised. Many real-world problems can be translated into optimisation ones, and the design of fuzzy models from the available data is no exception. Fuzzy system design and fuzzy model identification can be formulated as a search and optimisation problem in high-dimensional space, where each point corresponds to a fuzzy system, in other words, represents membership functions, rule-base, and hence the corresponding system behaviour. Given some objective/fitness function, the system performance forms a hyper-surface, and designing the optimal fuzzy system is equivalent to finding the optimal location on this hyper-surface. The hyper-surface is generally infinitely large, nondifferentiable, complex, noisy, multimodal, and deceptive (Shi, Eberhart, & Chen, 1999).

Figure 1. Fuzzy model identification using PSO

These characteristics make EAs better candidates for searching the hyper-surface than traditional, gradient-based methods. PSO algorithms, like GAs, are able to find optimal or near optimal solutions in a given complex search space, and can be used to modify/learn fuzzy model parameters. The idea of fuzzy model identification through PSO algorithm is illustrated in Figure 1.

Model Formulation

There are certain issues required to be addressed for solving any optimisation problem. These issues are:


- **Defining the solution space** — This involves defining the variables to be optimised and their respective domains.
- **Defining the constraints** — Here, it is required to define a set of constraints which must be followed by the candidate solutions.
- **Defining the fitness(objective) function** — The fitness/objective function not only represents the quality of each solution, but also acts as a link between the optimisation algorithm and the problem under consideration. It is imperative to select a good fitness function that accurately represents, in a single number, the goodness of the solution. Further, it is expected that the selected fitness function should exhibit a functional dependence that is relative to the importance of each characteristic being optimised.

The goal of optimisation is to find values that satisfy the defined constraints, and that maximise (or minimise) the fitness function.

Therefore, in order to use PSO for the identification of optimised fuzzy models successfully, we have to define the solution space, constraints, and the fitness function.

Another important consideration is the solution encoding, that is, how to represent a fuzzy system by a particle. In order for a particle to completely represent a fuzzy system,

Figure 2. A multi-input, single-output fuzzy model



all necessary information about the rule-base and membership functions is required to be specified. It is also advisable to evolve the membership functions and rule-base simultaneously, since they are co-dependent in a fuzzy system.

For the model formulation, consider a multi-input, single-output (MISO) system with n inputs as shown in Figure 2. The numbers of fuzzy sets for the inputs are $m_1, m_2, m_3, \dots, m_n$ respectively.

Some of the assumptions used for model formulation are:

1. Only triangular membership functions were used for both input and output variables.
2. The number of membership functions for each input and output variable were kept fixed.
3. The first and last membership functions of each input and output variable were represented with left- and right-skewed triangles.
4. The centres of all the triangular membership functions were fixed and were placed symmetrically over the universe of discourse.
5. A complete rule-base was considered (a rule-base is said to be complete when all possible combinations of input membership functions of all the input variables are considered).

It is important to mention here that a multi-input, multi-output (MIMO) model can be written as a *set* of MISO models.

Encoding Method for Membership Functions

Consider a triangular membership function, and let parameters x_k^l, x_k^c and x_k^r represent the coordinates of the left anchor, cortex, and right anchor of the k^{th} linguistic variable, as shown in Figure 3.


An efficient and convenient way to characterize this membership function is by using a parametric representation by means of the 3-tuple (x_k^l, x_k^c, x_k^r) .

Thus the parameters of the membership functions for the input and output variables are represented by the following particle:

$$(x_1^l, x_1^c, x_1^r, x_2^l, x_2^c, x_2^r, \dots, x_n^l, x_n^c, x_n^r, x_{n+1}^l, x_{n+1}^c, x_{n+1}^r) \quad (6)$$

The index $n+1$ corresponds to the membership functions of the output variable.

Figure 3. Characteristics of a triangular membership function



During the entire run of the PSO algorithm, the following constraints were imposed for every membership function of input and output variables:

$$x_k^l < x_k^c < x_k^r \quad (7)$$

At the same time, the overlapping between the adjacent membership functions is also ensured by imposing additional constraints. If for the purpose of simplicity we consider that a variable is represented by three fuzzy sets as in Figure 4, then those additional constraints to ensure overlapping can be represented as:

$$x_{min} \leq x_2^l < x_1^r < x_3^l < x_2^r \leq x_{max} \quad (8)$$

where x_{min} and x_{max} are the minimum and maximum values of the variable, respectively.

The additional constraints represented in Equation (8) can be generalized for any number of membership functions and are represented as:

$$x_{min} \leq x_2^l < x_1^r < x_3^l < x_2^r < \dots \dots \dots < x_{n-1}^l < x_{n-2}^r < x_n^l < x_{n-1}^r \leq x_{max} \quad (9)$$

The particle size required to encode the membership functions for each variable, while considering the assumptions made earlier in this section, can be represented as:

$$\text{Particle Size} = 2m_i - 2 \quad (10)$$


In Figure 4, $m_i=3$ and hence the particle size for encoding the variable consisting of three membership functions is $2*3-2=4$.

Equation (10) can be generalized for the fuzzy system of Figure 2. The dimensions of the particle to encode only membership functions for input and output variables for a Mamdani fuzzy model can be represented as in Equation (11).

$$\text{Particle size (for membership functions)} = \sum_{i=1}^{n+1} (2m_i - 2) \quad (11)$$

where n is the number of input variables, and m_i the number of fuzzy sets for i th input.

Figure 4. Representation of overlapping through constraints for a variable with three membership functions



As mentioned earlier, the index $n+1$ corresponds to the membership functions of the output variable.

Encoding Method for Fuzzy Rules

Let us again consider the system shown in Figure 2, taking into account the assumption that a complete rule-base is being considered. The particle size required for representing the entire rule-base is given by Equation (12).

$$\text{Particle size (for rule base)} = \prod_{i=1}^n m_i \quad (12)$$

Here each dimension represents the index of the membership functions of the output variable.

Particle dimensions required for encoding the Mamdani fuzzy model can be obtained in Equation (13) by simply adding Equations (11) and (12).

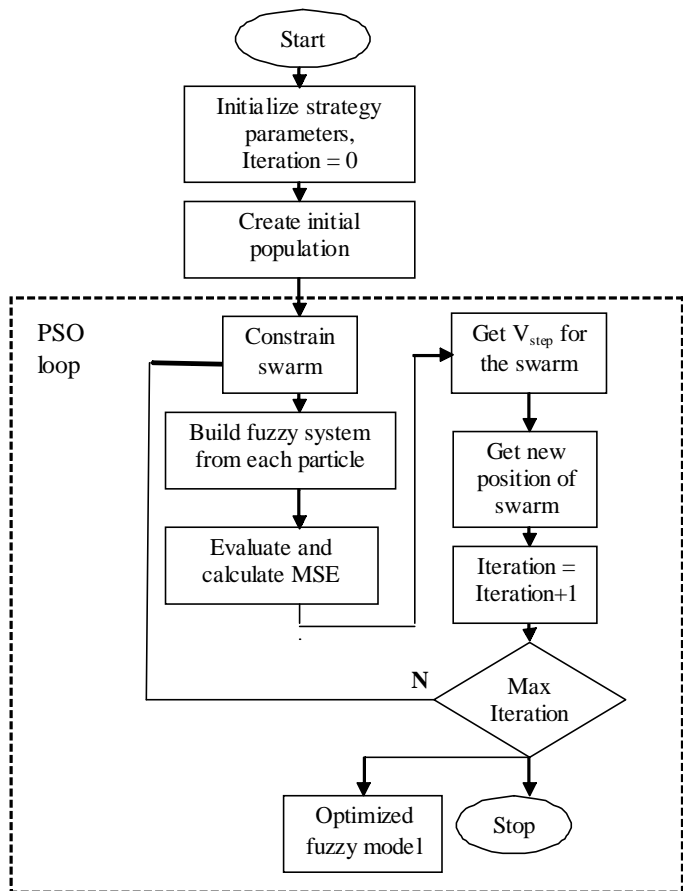
$$\text{Particle size (Mamdani model)} = \sum_{i=1}^{n+1} (2m_i - 2) + \prod_{i=1}^n m_i \quad (13)$$

If we consider the Singleton fuzzy model and assume that the consequent singleton values are t in number, then the particle dimensions required for encoding the model can be obtained simply from Equation (13) after a little modification, which is represented by Equation (14).

$$\text{Particle size (singleton model)} = \sum_{i=1}^n (2m_i - 2) + t + \prod_{i=1}^n m_i \quad (14)$$

The mean square error (MSE) defined in Equation (15) is used as the fitness function for rating the fuzzy model.

Figure 5. Flowchart of the framework for fuzzy model identification using PSO



$$MSE = \frac{1}{N} \sum_{k=1}^N [y(k) - \tilde{y}(k)]^2 \quad (15)$$

where $y(k)$ is the desired output and $\tilde{y}(k)$ the actual output, and N is the number of data points taken for model validation.

The framework for identifying the fuzzy model using the PSO algorithm is represented in Figure 5.

The corresponding pseudo-code for the used framework is listed as follows:

Begin

Define strategy parameters for PSO Algorithm;
 Iteration = 0;
 Create initial particle swarm;
while Iteration ≤ Maximum Iteration

Constrain Swarm;
Build fuzzy model for each particle;
Evaluate each fuzzy model & calculate MSE using Eqn.(15);
Get v_{id}^{n+1} as defined in Eqn.(1);
Determine new position of each particle by using Eqn.(2);
Iteration = Iteration+1;
end
End

Now it is possible that during the movement of the swarm, some particles may move out of the bounds defined by the system constraints. It is therefore necessary to constrain the exploration to remain inside the valid hyperspace. Whenever a particle moves to a point representing an invalid solution, it is reset within the valid bounds. Thus all the particles in the swarm are scrutinized after every iteration to ensure that they represent only valid solutions.

Application Example

Now let us consider the case of a Ni-Cd rapid battery charger, which is a two input, single output system. From an optimisation point of view, the input and output variables identified for the system have already been described in the previous section along with their universes of discourse.

Figure 6. Particle of 21 dimensions representing a Mamdani fuzzy model with two input variables, one output variable, and nine rules

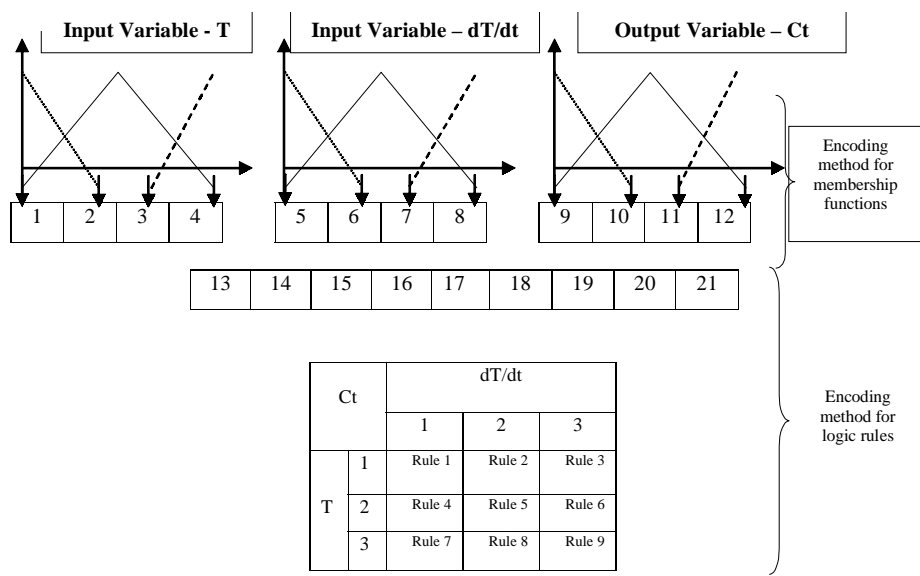
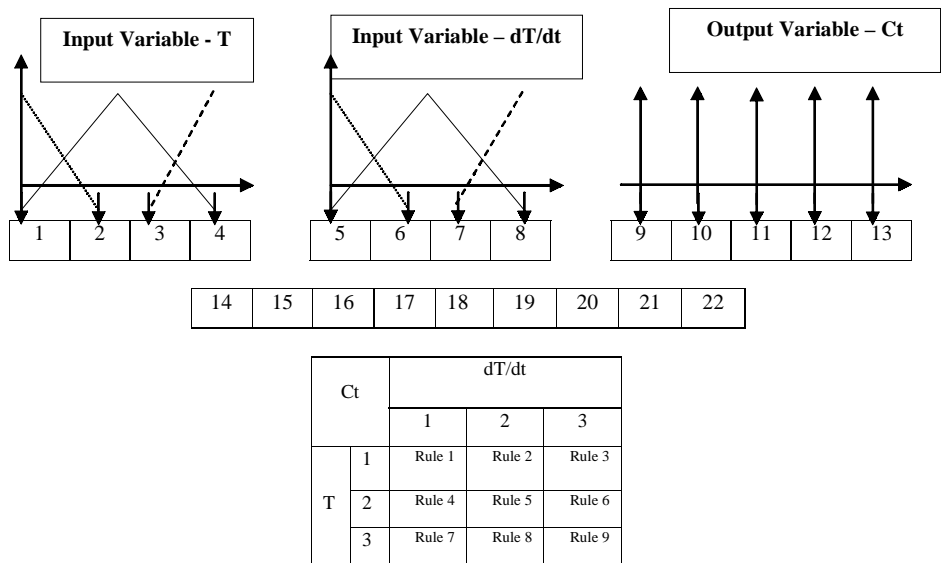


Figure 7. Particle of 22 dimensions representing a singleton fuzzy model with two input variables, one output variable, and nine rules



If we consider that each input and output variable for this system is represented by three fuzzy sets, then the particle size to encode a Mamdani fuzzy model may be calculated from Equation (13) as follows:

$$\sum_{i=1}^{n+1} (2m_i - 2) + \prod_{i=1}^n m_i \quad (13)$$

$$= (2 * 3 - 2) + (2 * 3 - 2) + (2 * 3 - 2) + 3 * 3 = 21$$

The particle representing a Mamdani fuzzy model for the Ni-Cd rapid battery charger is shown in Figure 6.

If we consider five possible consequent values for the singleton model, then the corresponding fuzzy model can be represented by a 22-dimension particle, obtained from Equation (14), as shown in Figure 7.

Simulation Results

The simulations were carried out using software tools developed by the authors, namely, the PSO toolbox (<http://sourceforge.net/projects/psotoolbox>) and PSO fuzzy modeler for Matlab (<http://sourceforge.net/projects/fuzzymodeler>). These tools are an open-source initiative and have been released under general public license (GPL). They are hosted on SourceForge.net, the world's largest open-source software development Web site.

Table 2. PSO algorithm parameters for fuzzy model identification in the Ni-Cd battery charger

Parameter	Value
Swarm size	30
Iterations	2500
c_1	2
c_2	2
w_{start} (inertia weight at the start of PSO run)	1
w_{end} (inertia weight at the end of PSO run)	0.1
V_{max}	100

Table 3. Simulation results

Model	MSE of fuzzy system corresponding to swarm's gbest		Simulation time
	After 1 st Iteration	After 2500 Iterations	
Mamdani	12.46	0.1455	19.424 hours
Singleton	46.94	0.1123	16.633 hours

The strategy parameters (for both the Mamdani and Singleton models) of the PSO algorithm used for fuzzy model identification of the Ni-Cd rapid battery charger are listed in Table 2, and the simulation results obtained for both models (obtained using an Intel Pentium-IV desktop computer with 256MB RAM) are presented in Table 3. Centre-of-Gravity and Weighted Average defuzzification techniques (Yen & Langari, 2003) were selected for the Mamdani and Singleton fuzzy models, respectively.

The results clearly show the effectiveness of the approach, as considerable improvement in the performance of fuzzy models was achieved after the complete run of the PSO algorithm. The longer simulation time for the Mamdani fuzzy model can be attributed to its more complicated defuzzification process.

The performance of the Singleton model identified through PSO was found to be even better than the Tagaki-Sugeno fuzzy model identified from the same data by using ANFIS model (Jang, 1993), where a MSE of 0.1321 was obtained (Khosla, Kumar, & Aggarwal, 2003b).

TAGUCHI METHODOLOGY AND EXPERIMENTS

In order to illustrate the effect of choice of strategy parameters, we conducted some experiments with different sets of parameters to identify the Mamdani fuzzy models for the Ni-Cd battery charger. The results are tabulated in Table 4, and the large variations

Table 4. Simulation results with different sets of strategy parameters (swarm size = 30; iterations = 2500)

c_1	c_2	w_{start}	w_{end}	V_{max}	MSE
0.5	0.5	0.9	0.1	50	10.229
1	1	0.9	0.4	100	6.898
1.5	1	2	0.3	50	0.3739
2	2	0.9	0.3	75	0.0489

in the values of fitness function (MSE) clearly indicate the dependence of the PSO algorithm performance on the choice of strategy parameters.

In this section, we shall discuss how, by using the Taguchi method, it is possible to arrive at good strategy parameters by performing a small number of experiments.

The main features of the Taguchi method are listed below (Bagchi, 1993; Chou, Chen, & Li, 2000; Ross, 1996; Taguchi, Chowdhury, & Wu, 2005; Tsai, Liu, & Chou, 2004):

- The fundamental principle of the Taguchi method, an important tool for robust design, is to improve the quality of a product by minimising the effect of the causes of variation, without eliminating the causes, per se.
- The two major tools used in the Taguchi method are the orthogonal array (OA) and the signal-to-noise-ratio (SNR).
- OA is a matrix in which the rows represent the level of factors in each run, and the columns represent a specific level that can be changed for each run. The OAs are represented by the following notation: $L_a(b^c)$, where a is the number of runs, b the number of levels, and c the number of columns. (In our experiments, we have used $L_{16}(4^5)$ to identify the best strategy parameters for PSO with five four-level factors).
- The array is referred to as “orthogonal” because all columns can be evaluated independently of one another.
- SNR is indicative of quality, and the purpose of the Taguchi experiment is to determine the best level for each operating parameter such that the SNR is maximised (or minimised).
- The OAs of the Taguchi method are *fractional* factorial designs that are used to study a large number of parameters with a small number of experiments. On the other hand, a full factorial design which represents the traditional or classical approach requires running *all* possible combinations (for example, consider a process that involves five four-level factors. The total number of experiments (combinations) = $4^5 = 1024$. If $L_{16}(4^5)$ OA is selected, only 16 experiments are required to be performed).

The steps of the Taguchi method (Ross, 1996; Taguchi, Chowdhury, & Wu, 2005) are depicted in the form of a flowchart in Figure 8 and are described as follows:

Step 0:	Start
Step 1:	In this step, it is required to define a clear and concise statement of the problem to be solved.
Step 2:	Implies identifying the objective function through some output measurable characteristics, which represent the quality characteristic.
Step 3:	Identify the control factors and their levels that shall influence the selected quality characteristic.
Step 4:	This step involves selection of an appropriate orthogonal array for the experiments.
Step 5:	Involves computer simulations (actual physical experiments) appropriate for the problem under consideration.
Step 6:	Generally, techniques like SN response graphs (Ross & Taguchi, 1996; Taguchi, Chowdhury & Wu, 2005) are used to analyze the data generated through experiments in Step 5.
Step 7:	The data analysis carried out in Step 6 helps in identifying optimum level of control factors.
Step 8:	A confirmation experiment is run with the optimum control factors obtained in Step 7. This is basically a validation or invalidation of optimum levels of control factors. An unsatisfactory confirmatory experiment implies that additional experiments are required to be performed (Go to Step 3).
Step 9:	Stop

The Taguchi method is divided into three main phases: (1) the planning phase, (2) the conducting phase, and (3) the analysis phase. The planning phase includes steps 1 through 4, step 5 is the conducting phase, and the analysis phase comprises steps 6 to 8.

Use of the Taguchi method for finding the optimum strategy parameters of the PSO algorithm for fuzzy model identification is elaborated below:

Step 0:	Start
Step 1:	The objective is to identify a good combination of strategy parameters so that the performance of the PSO algorithm for fuzzy model identification can be improved.
Step 2:	SNR represents the quality characteristic, and for the system under consideration is defined as: $SNR = 10/MSE$, where Mean Square Error (MSE) was defined earlier in Equation (15) (a high SNR implies good performance).
Step 3:	The factors (A-E) and the corresponding parameters are listed in Table 5. The levels of each operating parameter are listed in Table 6.
Steps 4 & 5:	OA $L_{16}(4^3)$, as discussed above and listed in Table 7, has been selected. The table also includes the results obtained through computer simulation. Some of the parameters fixed for the complete run are: Swarm size = 30; Iterations = 2500; Constriction factor = 1
Step 6:	From the results obtained from different experiments, a response table and response graphs have been obtained. The response table is used for recording the processed data, and it presents the calculations of effects from the orthogonally-designed experiments. The response graph is a graphical representation of the data presented in the response table, which can be used to quickly identify the effects of the different parameters.

Step 7:	From the response graphs, the optimum level of each factor can be predicted as the level that has the highest SNR value. Thus the optimal configuration identified from the response graphs is $A_3B_4C_2D_1E_2$.
Step 8:	The confirmatory experiment was run with the optimum control factors obtained in Step 7 and listed in Table 9. With these strategy parameters for PSO, the MSE obtained was 0.03679 and thus the SNR was 271.8130. The values of SNR obtained from all experiments in the selected OA, along with the confirmatory experiment with the optimum parameters, are plotted in Figure 10, which clearly shows that the SNR for the confirmatory experiment is indeed superior to all the SNR's listed in Table 7. Thus, the best performance was obtained with the parameters identified through the Taguchi method, which validates the approach.
Step 9:	Stop

From the data presented in Table 7, the entries for the response table (Table 8) are calculated as follows:

$$\text{Average SNR for } A_1 = \frac{0.9776 + 0.9797 + 32.0307 + 1.1496}{4} = 8.7844 \quad (18)$$

Figure 8. Flowchart of Taguchi optimisation methodology

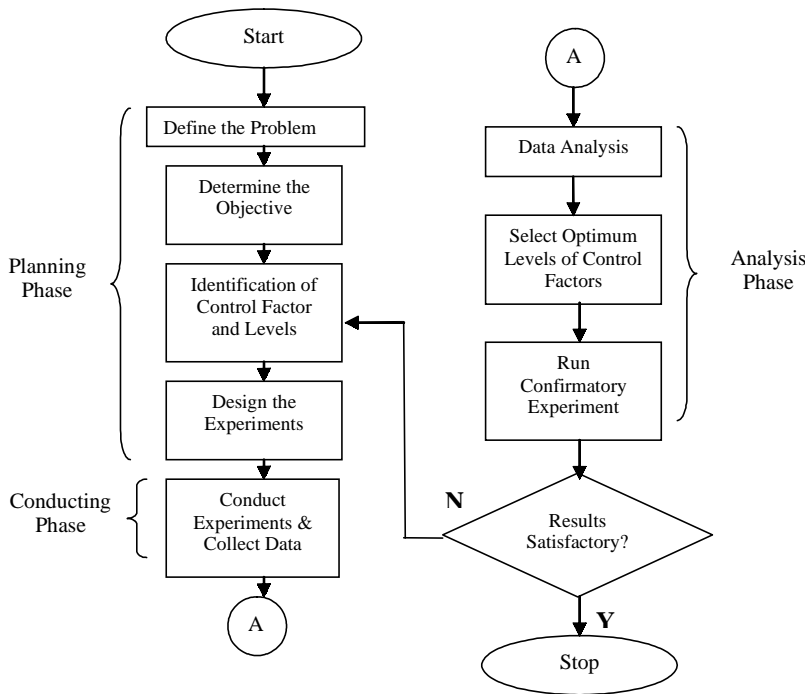


Table 5. Factors and their corresponding parameters

Factor	Corresponding strategy parameter of PSO
A	c_1
B	c_2
C	w_{start}
D	w_{end}
E	v_{max}

Table 6. Levels of the different parameters

Factor	Level			
	1	2	3	4
A	0.5	1	1.5	2
B	0.5	1	1.5	2
C	0.9	1	1.5	2
D	0.1	0.2	0.3	0.4
E	50	75	100	125

Table 7. $L_{16}(4^5)$ OA and results

Experiment number	Factor					MSE	SN ratio (10/MSE)
	A	B	C	D	E		
1	1	1	1	1	1	10.229	0.9776
2	1	2	2	2	2	10.2068	0.9797
3	1	3	3	3	3	0.3122	32.0307
4	1	4	4	4	4	8.6988	1.1496
5	2	1	2	3	4	9.8942	1.0107
6	2	2	1	4	3	6.8983	1.4496
7	2	3	4	1	2	0.1479	67.6133
8	2	4	3	2	1	3.584	2.7902
9	3	1	3	4	2	0.1105	90.4977
10	3	2	4	3	1	0.3739	26.7451
11	3	3	1	2	4	7.1799	1.3928
12	3	4	2	1	3	0.0389	257.0694
13	4	1	4	2	3	0.3359	29.7708
14	4	2	3	1	4	0.3094	32.3206
15	4	3	2	4	1	6.4665	1.5464
16	4	4	1	3	2	0.0489	204.4990

Table 8. Response table (SNRs)

Level	Factor				
	A	B	C	D	E
1	8.7844	30.5642	52.0798	89.4952	8.0148
2	18.2159	15.3737	65.1515	8.7334	90.8974
3	93.9263	25.6458	39.4098	66.0714	80.0801
4	67.0342	116.3770	31.3197	23.6609	8.9684

$$\text{Average SNR for } A_2 = \frac{1.0107 + 1.4496 + 67.6133 + 2.7902}{4} = 18.2159 \quad (19)$$

$$\text{Average SNR for } B_1 = \frac{0.9776 + 1.0107 + 90.4977 + 29.7708}{4} = 30.5642 \quad (20)$$

$$\text{Average SNR for } C_2 = \frac{0.9797 + 1.0107 + 257.0694 + 1.5464}{4} = 65.1515 \quad (21)$$

...and so on.

From the response table shown in Table 8, the response graphs can be derived (see Figure 9).

The predicted best strategy parameters are summarized in Table 9.

Figure 9. Response graphs





Table 9. Predicted best strategy parameters

Factor (Level)	Value
A(3)	1.5
B(4)	2
C(2)	1
D(1)	0.1
E(2)	75

Table 10. Computational overhead comparison

	Full factorial design (traditional)	Fractional factorial design (Taguchi method)
Time for 1 experiment	19.424 hours	19.424 hours
Total number of experiments (5 factors, each with 4 levels)	1024 (4^5)	16 (with $L_{16}(4^5)$ OA)
Total time for experimentation	828.16 days	12.94 days

Figure 10. SNR plots from all OA experiments and the confirmatory experiment



Comparison of Computational Overhead Between the Taguchi Method and the Traditional Approach

Assuming that the process of running the experiments is automated, and the experiments are being performed 24 hours a day, 365 days per year, the computational efforts required for the Taguchi method and for the traditional approach are compared in Table 10. The results show that huge savings in terms of number of experiments and hence the computation time can be achieved by following the Taguchi approach.

CONCLUDING REMARKS

In this chapter, we have described the use of the PSO algorithm for identification of optimised fuzzy models from the available data. Simulation results give a clear indication of the ability of PSO to arrive at good fuzzy models.

The suggested framework can be extended to increase the flexibility of the search by incorporating additional parameters so that the search for the optimal solution could be executed in terms of number of membership functions for each variable, the type of membership function, and the number of rules. Future work could be to investigate the influence of swarm size, and number of iterations, as well as possibly trying variants of the PSO algorithm for identifying fuzzy models with an objective to improve their performance further.

We have also presented the use of the Taguchi method to quickly identify the strategy parameters for the PSO algorithm. Computer simulations were carried out to confirm that improvements are achieved when the optimal operating parameters for the fuzzy model identification are obtained through the Taguchi method. Along the same lines, this approach can be used for finding a good set of operating parameters of PSO for any other system under consideration, with an objective to improve the performance. This can be achieved using fewer experiments, and hence little computational effort.

The proposed techniques are universal in nature, and there are no limitations to their usage. Future work could focus on using these methods for other fields and applications.

Throughout this chapter, we also introduced the *PSO toolbox* and *PSO fuzzy modeler for Matlab* open-source software tools. *PSO toolbox* is a collection of Matlab m-files that can be used to implement the PSO algorithm, whereas *PSO fuzzy modeler for Matlab* is capable of generating optimised fuzzy models from the available data.

REFERENCES

- Bagchi, T. P. (1993). *Taguchi methods explained: Practical steps to robust design*. New Delhi: Prentice Hall of India.
- Bastian, A. (1996, September 2-5). A genetic algorithm for tuning membership functions. *Proceedings of the 4th European Congress on Fuzzy and Intelligent Technologies — EUFIT(96)*, Aachen, Germany (Vol. 1, pp. 494-498).
- Buchmann, I. (1997, May 16-31). Getting the most out of your cell phone batteries. *Express Telecom*, 22-23.
- Carse, B., Fogarty, T. C., & Munro, A. (1996). Evolving fuzzy rule-based controllers using GA. *Fuzzy Sets and Systems*, 80, 273-294.
- Chou, J. -H., Chen, S. -H., & Li, J. -J. (2000). Application of the Taguchi-genetic method to design an optimal grey-fuzzy controller of a constant turning force system. *Journal of Material Processing Technology*, 105, 333-343.
- Clerc, M. (1999, July 6-9). The swarm and the queen: Towards a deterministic and adaptive particle swarm optimization. *Proceedings of the Congress on Evolutionary Computation*, Washington, DC (pp. 1951-1957).
- Eberhart, R. C., & Kennedy, J. (1995, October 4-6). A new optimiser using particle swarm theory. *Proceedings of the 6th Symposium on Micro Machine and Human Science*, Nagoya, Japan (pp. 39-43).

- Eberhart, R. C., & Shi, Y. (2001, May 27-30). Particle swarm optimisation: Developments, applications, and resources. *Proceedings of the IEEE Congress on Evolutionary Computation*, Seoul, South Korea (pp. 81-86).
- Hellendoorn, H., & Driankov, D. (Eds.). (1997). *Fuzzy model identification: Selected approaches*. Berlin: Springer-Verlag.
- Jang, J.-S.R. (1993). ANFIS: Adaptive-network-based fuzzy inference system. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(3), 665-685.
- Kennedy, J., & Eberhart, R. (1995). Particle swarm optimisation. *Proceedings of the IEEE Conference on Neural Networks*, Perth, Australia (Vol. IV, pp. 1942-1948).
- Kennedy, J., & Eberhart, R. (2001). *Swarm intelligence*. San Francisco: Morgan Kaufmann.
- Khosla, A. (1997). *Design and development of RFC-10: A fuzzy logic based rapid battery charger for nickel-cadmium batteries*. Master of technology thesis, Department of Electronics, Communication and Computer Engineering, National Institute of Technology (Formerly Regional Engineering College) Kurukshetra/ Kurukshetra University, Kurukshetra, India.
- Khosla, A., Kumar, S., & Aggarwal, K. K. (2002, December 18-21). Design and development of RFC-10: A fuzzy logic based rapid battery charger for nickel-cadmium batteries. *Proceedings of the High-Performance Computing Workshop on Soft Computing — HiPC2002*, Bangalore, India (pp. 9-14).
- Khosla, A., Kumar, S., & Aggarwal, K. K. (2003a, July 24-26). Identification of fuzzy controller for rapid nickel-cadmium batteries charger through fuzzy c-means clustering algorithm. *Proceedings of the North American Fuzzy Information Society — NAFIPS — Conference*, Chicago (pp. 536-539).
- Khosla, A., Kumar, S., & Aggarwal, K. K. (2003b, July 24-26). Fuzzy controller for rapid nickel-cadmium batteries charger through adaptive neuro-fuzzy inference system (ANFIS) architecture. *Proceedings of the North American Fuzzy Information Society — NAFIPS — Conference*, Chicago (pp. 540-544).
- Linden, D. (Ed.). (1995). *Handbook of batteries*. New York: McGraw Hill.
- Melin, P., & Castillo, O. (2005). Intelligent control of a stepping motor drive using an adaptive neuro-fuzzy inference system. *Information Sciences*, 170(2-4), 133-151.
- Nelles, O. (1996, September 2-5). FUREGA-Fuzzy rule extraction by GA. *Proceedings of the 4th European Congress on Fuzzy and Intelligent Technologies - EUFIT(96)*, Aachen, Germany (Vol. 1, pp. 489-493).
- Nozaki, K., Morisawa, T., & Ishibuchi, H. (1996). Adjusting membership functions in fuzzy rule-based classification systems. *Proceedings of the 3rd European Congress on Fuzzy and Intelligent Technologies — EUFIT(95)*, Aachen, Germany (Vol. 1, pp. 615-619).
- Parsopoulos, K. E., & Vrahatis, M. N. (2002). Recent approaches to global optimization problems through particle swarm optimization. *Natural Computing*, 235-306.
- Ross, P.J. (1996). *Taguchi techniques for quality engineering*. New York: McGraw Hill.
- Setnes, M., & Roubos, J. A. (1999, June 10-12). Transparent fuzzy modelling using clustering and GAs. *Proceedings of the North American Fuzzy Information Society — NAFIPS — Conference*, New York (pp. 198-202).
- Shi, Y., & Eberhart, R. (2001). Fuzzy adaptive particle swarm optimization. *Proceedings of the IEEE International Conference on Evolutionary Computation* (pp. 101-106).

- Shi, Y., Eberhart, R., & Chen, Y. (1999). Implementation of evolutionary fuzzy systems. *IEEE Transactions on Fuzzy Systems*, 7(2), 109-119.
- Taguchi, G., Chowdhury, S., & Wu, Y. (2005). *Taguchi's quality engineering handbook*. New York: John Wiley & Sons.
- Tsai, J. -T., Liu, T. -K., & Chou, J. -H. (2004). Hybrid Taguchi-genetic algorithm for global numerical optimization. *IEEE Transactions on Evolutionary Computation*, 8, 365-377.
- Wan, S. (1996). *The chemistry of rechargeable batteries*. Richmond, BC, Canada: Cadex Electronics Inc. (Revised August 26).
- Xie, X. -F., Zhang, W. -J., & Yang, Z. -L. (2002). Adaptive particle swarm optimisation on individual level. *Proceedings of the International Conference on Signal Processing — ICSP2002*, Beijing, China (pp. 1215-1218).
- Yen, J., & Langari, R. (2003). *Fuzzy logic: Intelligence, control and information* (3rd ed.). Upper Saddle River, NJ: Prentice Hall.

About the Authors

John Fulcher is currently a professor of information technology in the School of IT and Computer Science and director of the Health Informatics Research Centre at the University of Wollongong, Australia. He holds a BEE with Honours from the University of Queensland (1972), a Research Masters from LaTrobe University, Melbourne (1981), and a PhD from the University of Wollongong (1999). John is a member of the Association for Computing Machinery and a senior member of the Institute of Electrical and Electronic Engineers. Apart from serving as a preferred reviewer for ACM Computing Reviews, he is on the editorial boards of both *Computer Science Education* and *Intl. J. Intelligent & Fuzzy Systems*, and serves as reviewer for 13 other journals (including *IEEE Transactions*). His 100 or so publications include a best-selling textbook on microcomputer interfacing, a recent Springer research monograph on applied intelligent systems (co-edited with Prof. Lakhmi Jain), and three book chapters in the Oxford University Press Handbook of Neural Computing. Professor Fulcher was an invited keynote speaker at the 5th National Thai Computer Science and Engineering Conference. His research interests include microcomputer interfacing, computer science education, artificial neural networks (especially higher-order ANNs), health informatics, and parallel computing.

* * * * *

Ajith Abraham is currently an IITA distinguished visiting professor at Chung-Ang University, Korea. He holds a PhD from Monash University, Australia, and an MS from Nanyang Technological University, Singapore. Before becoming a full-time academic (for the past six years), he had been working with three multi-national companies on several industrial research and development projects over a seven-year period. His primary research interest is computational intelligence, with application to information

security, bio-informatics, web intelligence, energy management, financial modelling, and so forth. He has co-authored over 100 research publications in peer-reviewed journals and international conference proceedings, four of which have won “best paper” awards/nominations. More information can be found at <http://ajith.softcomputing.net>.

K. K. Aggarwal is the founder and vice-chancellor of Guru Gobind Singh Indraprastha University, Delhi, India. He has been the chairman of the Department of Electronics, Communication, and Computer Engineering at the National Institute of Technology (formerly, Regional Engineering College), Kurukshetra. He has also worked as an academic dean in this institute, as well as director of the Centre for Excellence for Manpower Development in Reliability Engineering, established by the Indian Ministry of Human Resource Development. Before taking up his present assignment, he was pro-vice chancellor at Guru Jambheshwar University, Hisar. He has worked extensively in various areas of electronics and computer engineering, publishing about 200 papers in reputed journals. He has been invited to deliver lectures in several universities in India and abroad, such as the University of Berkley, University of Cincinnati, and Florida State University in the USA; the University of Birmingham in the UK; and the Technical University of Germany. Besides his contribution in the academic field, he has made a very strong impact on the industrial world. He has been widely consulted by industry, delivered lectures in several industrial and research organizations, and conducted courses for many important industrial houses, such as Tata Consultancy Services; Bharat Heavy Electricals, Ltd.; TVS Electronics; DCM Data Systems; the Central Mechanical Engineering Research Institute; the Electronics Corporation of India; Bharat Electronics, Ltd.; the Indian Space Research Organization and Defence Research Laboratories, and similar. His most-notable contribution has been towards the reliability analysis for the PSLV (Polar Satellite Launch Vehicle). He was honoured by the Reliability Society of IEEE, USA, for his services as guest editor for a special issue on “State of Reliability Effort of the Indian Sub-Continent”. He was conferred the L. C. Verman Award by the Institute of Electronics and Tele-Communication Engineers, India. He received the Indra Gandhi “Priyadarshini” Award for his outstanding service, achievement in, and contribution to the field of education.

Quan Bai is a PhD candidate in computer science at the University of Wollongong, Australia, investigating coordination mechanisms for self-interested multi-agent systems. He received his double Bachelors degrees in computer science and international business from Tianjin University, China, in 2001, and his Master of Computer Science from the University of Wollongong, Australia, in 2002. He was a research assistant for an Australian Research Council Discovery project in 2003, and to date has published eight research papers.

Nick Barnes received his BSc and PhD from the Department of Computer Science and Software Engineering at the University of Melbourne. Since October, 2003, he has been a researcher with the Autonomous Systems and Sensing Technologies Programme at the National Information and Communications Technology Australia (NICTA). Prior to this, he was a lecturer in the Department of Computer Science and Software Engineering at the University of Melbourne for more than three years, as well as spending six months as a post-doctoral fellow with the LIRA-Lab at the University of Genoa, Italy. He also spent

two years with Accenture, working in the telecommunications industry. His research interests focus on perception and action, principally in developing effective perceptual systems to facilitate the action of robots, as well as for diagnosis with medical image analysis. Nick also has a strong interest in general robotics and perception in biology.

Russell Beale leads the Advanced Interaction Group in the School of Computer Science at the University of Birmingham, UK, a loosely-coupled cluster of people interested in human-computer interaction and related issues. His research focus is on using intelligence to support user interaction. This takes two forms: the use of artificial intelligence techniques in interactive systems, and better design based on analysing and modelling the user and their interactions with the system. Coupling AI with a clearer understanding of the user and their skills and limitations allows us to design more effective, usable systems.

Tony Farmer currently leads the Intelligent Systems research group within the CSIRO Division of Industrial Physics in Sydney, Australia. He holds a BSc with Honours (1965) and a PhD (1969) in physics from the University of Adelaide, South Australia. Dr. Farmer's research interests have covered spectroscopic, optical, electromagnetic, acoustic, and ultrasonic sensing, and the application of plasma technologies to industrial processes. More recently, he has been involved with investigations into the concepts of large sensing networks and their potential to make intelligent decisions based on the sensed data.

Luke Fletcher is a doctoral candidate under the supervision of Professor Alex Zelinsky at the Department of Information Engineering, Research School of Information Sciences and Engineering, Australian National University, Australia. In 1996, Luke completed his BE (with Honours) and BSc degrees at the University of Melbourne. His research interests include robust and real-time computer vision, stereo, optic flow, visual geometry, perception, and visual interfaces.

Russell Gluck is a lecturer in learning development at the University of Wollongong, Australia. He has worked for more than 20 years in squatter camps, prisons, and the Aboriginal Education Centre at the University of Wollongong. This wealth of experience has led to the development of an engagement process that enables literacy-inefficient people to draw, compose, write, and read stories, sufficient in some cases to successfully complete university studies.

Markus Hagenbuchner was born in Germany, and received his Bachelor in Computer Science (with Honours) from the University of Wollongong, Australia, in 1997. He worked as a research assistant in the Department of Artificial Neural Computing at the University of Ulm, Germany, from 1997 to 1998. In 1998, he returned to the University of Wollongong to commence a PhD in the Faculty of Informatics, which he subsequently received in 2002. Markus is currently working as a senior research associate at the University of Wollongong.

Mark Hedley is a principal research engineer at the Commonwealth Scientific and Industrial Research Organisation (CSIRO) ICT Centre, Sydney, Australia. Prior to this,

he was a lecturer at the University of Sydney. He received his BSc (1987) degree in computer science and pure mathematics, his BE (1989) degree with University Medal in electronic engineering, and his PhD (1992) degree in medical imaging, all from the University of Sydney. Dr. Hedley is a member of the IEEE, and the author of over 70 research publications in areas including sensor networks, communications theory, signal processing, image processing, and medical imaging.

Nigel Hoshcke is an electrical engineer working at the Commonwealth Scientific and Industrial Research Organisation (CSIRO) Division of Industrial Physics, Australia. In his current position, he is responsible for the development and implementation of multi-agent algorithms on a de-centralised network of digital signal processors (DSP) and sensors. Mr. Hoshcke graduated from the University of New South Wales in 2003 with a BE in electrical technology with first-class honours. His research interests include multi-agent systems, signal and image processing, and efficient implementation of DSP algorithms.

Lakshmi Jain is a director/founder of the Knowledge-Based Intelligent Engineering Systems (KES) Centre, located in the University of South Australia, Australia. He is a fellow of the Institution of Engineers, Australia. He has initiated a post-graduate stream by research in the knowledge-based intelligent engineering systems area. Professor Jain has presented a number of keynote addresses in international conferences on knowledge-based systems, neural networks, fuzzy systems, and hybrid systems. He is the founding editor of the *International Journal of Knowledge-Based Intelligent Engineering Systems* and *International Journal of Hybrid Intelligent Systems*, and served as an associate editor of the *IEEE Transactions on Industrial Electronics*. Professor Jain also initiated the First International Conference on Knowledge-Based Intelligent Electronic Systems in 1997, which has since become an annual event. He served as vice president of the Electronics Association of South Australia during 1997. Professor Jain's research interests focus on the application of novel techniques such as knowledge-based systems, artificial neural networks, fuzzy systems, intelligent agents, and evolutionary computation.

Arun Khosla is currently working as a senior lecturer in the Department of Electronics and Communication Engineering, National Institute of Technology (formerly, Regional Engineering College), Jalandhar, India. His research areas of interest are soft and biologically-inspired computing. He has been the co-ordinator of various Government of India-funded projects. As one of the founder members of the Department of Electronics and Communication Engineering, he has been associated with the development of various facilities. He has organised a number of training programs/tutorials in the field of soft computing at different institutes. He is a life member of the Indian Society of Technical Education.

Shakti Kumar is currently a professor and additional director at Haryana Engineering College (HEC), Jagadhari, India. He is the founding director of the Centre for Advanced Technologies in the same college. He has also served as a faculty member in the Department of Electronics, Communication, and Computer Engineering, National Institute of Technology (formerly, Regional Engineering College), Kurukshetra, the Depart-

ment of Computer Engineering, Atlim University, Ankara, Turkey, and the BITS Pilani-Dubai Center, UAE. He is a member of various professional bodies, such as IEEE, IETE, and ISTE. He has been associated with program and syllabi development for Kurukshetra University, Guru Jambheshwar University, Hissar, GGS Indraprastha University, Delhi, SLIET Longowal, Agra University, Agra, Atlim University (Turkey), and BITS Pilani. Before joining HEC, he held various positions, including the coordinator of the Centre for Excellence in Reliability Engineering at REC Kurukshetra, director at the Jind Institute of Engineering and Technology, Jind, and the Academy of Business and Engineering Sciences, Ghaziabad. He was a member of the World Scientific Committee for the International Conference held in 2004 in Greece, and the National Conference on Rural Development held at Agra University, Agra. He has published more than 75 research papers in international/national journals and conferences. He was awarded the “Rashtriya Nirman” Award by the International Business Council, New Delhi, India.

Chris Lewis is currently a senior technical officer with the Commonwealth Scientific and Industrial Organisation (CSIRO) Division of Industrial Physics, Sydney, Australia. He received a Certificate in electronics and communications (1982) and Diploma in electrical technology (1998) from the North Sydney Institute of Technical and Further Education. He is currently completing a BE in electrical technology at the University of Technology, Sydney. Mr. Lewis has worked as a technician with Medtel P/L (a medical technology distributor) servicing surgical lasers, gamma cameras, ultrasound imagers, and training hospital staff in their operation. He joined CSIRO in 1996 and has led the security radar project since 1999. Mr. Lewis has designed electronic systems for several projects within CSIRO, including squid magnetometers, ultra-wideband radars, and distributed computing hardware.

Rinku Panchal received her BIT with Honours degree from Griffith University. She received a PhD scholarship and started her PhD at Griffith University in 2003. In 2004, she moved to Central Queensland University, Australia, where she is currently a PhD candidate in the Faculty of Informatics and Communication. Her research interests include computational intelligence and pattern recognition. She has published four papers in the area of neural networks and pattern recognition.

Lars Petersson received his MSc in engineering physics (1997) and his PhD (2002) from the Royal Institute of Technology in Stockholm, Sweden. Prior to joining the Autonomous Systems and Sensing Technologies Programme at the National Information and Communications Technology Australia (NICTA), Australia, he was a researcher at the Australian National University in Canberra. His research interests are mainly in the area of active perception and action, and how this can be applied to, for example, mobile robotics, mobile manipulation, and most recently driver assistance systems.

Geoff Poulton is a senior principal research scientist within the CSIRO ICT Centre, Sydney, Australia. He has had a long research career within universities and CSIRO, with his main areas of expertise being originally in electromagnetic theory and antennas, and latterly in intelligent systems, complex systems, and biometrics. His most recent research interest has been in the area of the design and control of complex systems — a difficult

task, since such systems exist on the “edge of chaos”. He has produced over 130 publications and patents, has been awarded two CSIRO Medals, and has a long-standing interest in cross-disciplinary research. Dr. Poulton is also the director of the Australian Biometrics Institute.

Don Price received the BSc with Honours and PhD degrees in Physics from Monash University, Australia, in 1966 and 1970, respectively. Since joining CSIRO in 1982, he has been involved in the development of ultrasonic transducers, studies of the physics of linear and non-linear elastic wave propagation in inhomogeneous materials, with particular emphasis in recent years on propagation in planar-layered structures, and the development of ultrasonic measurement techniques based on these studies. He is currently the divisional science leader for CSIRO Industrial Physics, Australia, and is the project leader of the Ageless Aerospace Vehicles project, a collaborative investigation supported by CSIRO and the Non-Destructive Evaluation Sciences Branch of NASA at Langley Research Center, aimed at integrated smart sensing and communication systems.

Mikhail Prokopenko is a senior research scientist at the Commonwealth Scientific and Industrial Research Organisation (CSIRO) ICT Centre, Autonomous Systems Laboratory, Sydney, Australia. He holds a PhD in computer science from Macquarie University, Sydney, an MA in economics from the University of Missouri-Columbia, USA, and an MS in applied mathematics from the Azerbaijan Oil and Chemistry Institute, USSR. His long-term research goal is the design and control of complex self-organising multi-agent systems, and an effective integration of real-time perception, reasoning, and action in autonomous agents. His work has resulted in over 50 publications and patents. In June, 2002, Dr. Prokopenko received the Japanese Society for Artificial Intelligence award for scientific contribution to the RoboCup Simulation League. Currently he is coordinating the Multi-Agent Algorithms team within the CSIRO-NASA Ageless Aerospace Vehicle research project.

Andy Pryke received his PhD in data mining and visualization from the University of Birmingham, UK, in 1998. Since then he has pursued dual careers in academia and business, gaining considerable experience as a project leader and architect in commercial software development projects. His research interests include evolutionary computation, interactive visualization, and data mining. Andy currently works for the Centre of Excellence for Research in Computational Intelligence and Applications, a technology transfer unit at the University of Birmingham which is dedicated to using nature-inspired computation techniques to solve real business problems.

Andrew Scott graduated with a BSc with first-class Honours in 1980 and a PhD in physics in 1984 from Macquarie University, Sydney, Australia. He has worked on the physics of experimental lasers, atomic and molecular spectroscopy, corona discharges, the physics of welding arcs, plasma processing, and the interaction of laser radiation with materials for industrial applications. Dr. Scott is currently active in the field of ultrasonics for the characterization of materials, and more generally in the development of “intelligent” sensing systems. He is presently a principal research scientist at CSIRO Industrial Physics, Sydney.

Amanda J. C. Sharkey has an inter-disciplinary background. After taking a first degree in psychology, she held a variety of research positions at the MRC Cognitive Development Unit, University of Exeter, UK, then Yale University and Stanford University in the USA. She completed her PhD in psycho-linguistics in 1989 at the University of Essex. Since then Dr. Sharkey has conducted research in neural computing at the University of Exeter, before moving to the University of Sheffield, UK, where she is now a senior lecturer in the Department of Computer Science. Her other research interests include combining neural nets and other estimators, and more recently, robotics. Amanda has over 50 publications to her credit. She is a founding member of the scientific committee for the international series of workshops on multiple classifier systems, and is an editor of the journal *Connection Science*.

Noel Sharkey is a professor in the Department of Computer Science at the University of Sheffield, UK, and EPSRC senior media fellow (2004-2007). He has held a number of research and teaching positions in the UK (Essex, Exeter, and Sheffield) and in the USA (Yale, Stanford, and Berkeley). Noel has moved freely across academic disciplines, lecturing in departments of engineering, philosophy, psychology, cognitive science, linguistics, artificial intelligence, and computer science. He holds a Doctorate in experimental psychology, is a chartered electrical engineer, and a chartered information engineer. Professor Sharkey has published over 100 academic articles and books, as well as articles and web chats for BBC web pages and regular magazine articles. In addition to editing several journal special issues on modern robotics, Noel is editor-in-chief of *Connection Science*, and an editor of both *Robotics and Autonomous Systems* and *Artificial Intelligence Review*. Professor Sharkey appears regularly on TV (over 250 appearances in the last seven years) and is interviewed regularly on radio, in magazines, and in newspapers. He has been chief judge for every series of *Robot Wars* throughout the world, as well as “techspert” for four series of *TechnoGames*.

Cong Tran currently works in the Knowledge Engineering System Group, School of Electrical and Information Engineering of the University of South Australia, Australia. He received his PhD, ME, Graduate Diploma, and BE degrees from the University of South Australia. His research interests include artificial intelligence, machine learning, knowledge acquisition, and hybrid intelligent techniques. Application areas include data mining, industrial intelligent control, and decision support systems.

Phuong Kim To was born in Danang City, Vietnam, in 1983. She attended Nguyen Khuyen Gifted Secondary School, and Le Quy Don Gifted High School. Phuong Kim received a Bachelor of Computer Science from the University of Wollongong in 2004. While studying at the University, she worked part-time as a tutor and research assistant in the Faculty of Informatics. In 2004, Phuong Kim took a position as a software developer for Tedis P/L, Australia.

Ah Chung Tsoi was born in Hong Kong. He received the higher diploma in electronic engineering from Hong Kong Technical College in 1969, the MSc degree in electronic control engineering, and a PhD in control engineering from the University of Salford in 1970 and 1972, respectively. He also received a BD from the University of Otago, New Zealand, in 1980. From 1972 to 1974, he was a senior research fellow at the Inter-University

Institute of Engineering Control, University College of North Wales, Bangor. From 1974 to 1977, he was a lecturer at the Paisley College of Technology, Scotland. From 1977 to 1985, he was a senior lecturer at the University of Auckland, New Zealand, and between 1985 and 1990, he was a senior lecturer at the University College, University of New South Wales. From 1990 to 1996, Ah Chung Tsoi was first associate professor, then professor in electrical engineering at the University of Queensland, Australia, during which time he led a large research and development project on speaker verification. Professor Tsoi spent the years 1996 to 2004 at the University of Wollongong, Australia, first as the dean of informatics (1996-2001), then director of information technology services (1999-2001), and pro-vice chancellor for ITC (2001-2004). Professor Tsoi initiated a number of campus infrastructure projects in these latter two roles, including gigabit Ethernet, voice-over IP telephony, wireless access, and a student management package. His research interests include aspects of neural networks and fuzzy systems, and their application to practical problems, adaptive signal processing, speech processing, and adaptive control. He applies these techniques to many practical situations, including data mining and internet search engine design. In 2004, Professor Tsoi took the position of executive director of mathematics, Information and Communication Sciences, Australian Research Council, and in 2006 to Monash University in Melbourne, Australia, to become director of their e-Research Centre. He publishes widely in the area of adaptive signal processing, neural networks, and fuzzy systems. In 2001, Professor Tsoi received the Outstanding Alumni Award from the Hong Kong Polytechnic University.

Philip Valencia received his BE in electronics and BIT from the Queensland University of Technology in 2000. Joining the CSIRO, Australia, in 2001, Philip has gained extensive experience in complex systems, image processing and active vision, machine learning, and artificial intelligence. His recent work has taken him to focus on robotics in heterogeneous sensor networks, including actuation, mobile robots in self-managing ad hoc sensor networks, and autonomous robotic repair of a multi-cellular ageless aerospace vehicle surface. Philip is presently facilitating CSIRO-wide interactions on engineering, design, and control of complex systems.

Brijesh Verma received his PhD degree from the Warsaw University of Technology, Poland. Currently, he is an associate professor in the Faculty of Informatics and Communication at Central Queensland University, Australia. His research interests include computational intelligence and pattern recognition. He has published over 80 papers, one book, two book chapters, and eight edited books (conference proceedings). He is a recipient of 10 research grants from national/international organizations and industry. He has supervised 26 research students in his area of research. He is a co-editor in chief of the *International Journal of Computational Intelligence and Applications*, an associate editor of *IEEE Transactions on Biomedicine in Information Technology* and an editorial board member of the *International Journal of Hybrid Intelligent Systems*. He has served on the program/organizing committees of over 15 conferences. He has been listed in the 20th Edition of *Marquis Who's Who in the World*. He is a member of the IEEE Computational Intelligence Society.

Peter Wang graduated with a BSc in computer science and information systems from Macquarie University, Sydney, in 2000. He is currently working in the Commonwealth

Scientific and Industrial Research Organisation (CSIRO) ICT Centre, Autonomous Systems Laboratory, Sydney, Australia, on the research and development of autonomous agents and multi-agent systems. He has designed and developed multi-agent systems in a number of application areas, including RoboCup simulation and reconfigurable sensor and communication networks (the CSIRO-NASA Ageless Aerospace Vehicle project). In June, 2002, Peter received the Japanese Society for Artificial Intelligence award for scientific contribution to the RoboCup Simulation League, for work on the entropy of joint beliefs as a measure of multi-agent coordination potential.

Alexander Zelinsky is currently the director of the Commonwealth Scientific and Industrial Research Organisation (CSIRO) ICT Centre, Sydney, Australia. Prior to this, he was a professor in the Research School of Information Sciences and Engineering at the Australian National University, Canberra. He received his Bachelors (1983) and PhD (1991) degrees in computer science and electrical engineering from the University of Wollongong, Australia. Professor Zelinsky has worked as a research scientist in Japan (1992-1995) at the University of Tsukuba and the Electro Technological Laboratory. In 1996, he joined the ANU to lead the Robotics Systems Laboratory. His primary research interests are in mobile robotics, multi-robot systems, human-robot interaction, and real-time computer vision systems. In recent years, the focus of his work has moved to building systems that are suitable for reliable deployment in the real world, with the ultimate goal of developing robotic technology for mass market applications. To this end, Professor Zelinsky founded Seeing Machines P/L, a company that is building machine vision applications for the automotive industry.

Minjie Zhang is currently an associate professor in computer science and software engineering at the University of Wollongong, Australia. She received a BSc in computer science from Fudan University, China, in 1982, and a PhD from the University of New England, Australia, in 1996. She is a member of IEEE, the IEEE Computer Society, and ISCA. She is the author or co-author of over 60 research papers. Her research interests include distributed artificial intelligence, distributed information retrieval, agent-based meta-computing, and agent-based software engineering.

Index

A

acoustic features 73
 active damage evaluation (ADE) 192
 active pedestrian protection systems (PPS) 134
 adaptive dead reckoning scheme (ADRS) 217, 220
 adaptive network-based fuzzy inference system (ANFIS) 10
 aerospace vehicle 186, 196
 ageless aerospace vehicle (AAV) 196
 agent 52, 159
 agent communication languages (ACLs) 53
 agent interaction 52, 57
 agent interaction protocol 52, 58
 ALLIANCE architecture 165
 ant colonies 159, 204, 217
 ant colony optimisation (ACO) 187, 217, 220
 ant colony optimisation (ACO) algorithm 187
 architecture for behaviour-based agents, (ABBA) 162
 artificial neural network (ANN) 1, 6, 78, 162, 251, 258, 262

artificial neural network classifier 258, 262
 artificial neural network learning 1
 auto-associator 251, 253, 257, 262, 268
 auto-associator neural network (AANN) 264, 267
 auto-regressive modeling 74
 automatic speech recognition (ASR) 69
 autonomous robots 163

B

back-propagation (BP) 78, 256
 back-propagation neural networks (BPNN) 256, 267
 bag of words 33
 battery charger 227, 284
 biological inspiration 160
 blind spot monitoring 143
 breast cancer 252
 breast imaging reporting and data system (BI-RADS) 257

C

C4.5 243
 cell-replication 214
 cells 196
 chain code 253, 256

classification 129, 258
 classification and regression trees (CART)
 9, 21
 classifier 258, 262
 closed impact boundary 206
 cognitive task analysis (CTA) 5
 cognitive work analysis (CWA) 5
 collective robotics 158
 collision mitigation brake system (CMS)
 123
 colour segmentation 130
 coloured Petri net (CPN) 52, 54, 57
 combat air patrol (CAP) 4
 ComDot™ 84
 communications 194
 computer vision 135
 concept demonstrator (CD) 189, 229
 concept demonstrator (CD) system 189
 condition-based maintenance (CBM) 195
 Confirmed Caller™ 83
 conjugate gradient algorithm (CGA) 7
 connected boundary-fragment (CBF) 209
 connected trail-fragment (CTF) 222
 cooperative tasks 172
 cranio caudal view (CC) 255
 critical damage 192, 202
 cross correlation 131

D

damage 187
 damage scenario 202
 DAML+OIL 55
 DAMN architecture 163
 dandelion head 240
 data identification 242
 data mining 243, 245, 249
 data value constraints 60
 data visualisation 241
 datasets 245
 dead reckoning scheme (DRS) 220
 decentralised control 163
 decision factor (DF) 3
 decision problem (DP) 3
 decision rules 5
 decision support system (DSS) 1
 decision trees 2, 6

design parameters 213
 detection algorithm 253
 digital database for screening mammogra-
 phy (DDSM) 255
 digital mammogram 252
 direct solution method (DSM) 258
 direct solution method-based (DSMNN)
 267
 disparity map 147
 distillation algorithm 124
 draw-talk-write (DTW) 68, 84, 110
 driver assistance systems (DAS) 120
 dynamic time warping 75

E

evolutionary algorithm-based (EANN) 267
 evolutionary algorithms (EA) 251, 260,
 evolutionary computation (EC) 2
 evolutionary neural network 251, 260
 evolutionary programming and probabilis-
 tic (EPPNN) 267
 exploration 173

F

fast radial symmetry detector 130
 feature extraction 253, 256
 feature selection 253, 258, 260
 feed-forward network 256
 fitness function 227
 frame boundary 206
 fuzzy inference method 1
 fuzzy inference system (FIS) 7
 fuzzy logic 2
 fuzzy model identification 276, 278
 fuzzy neural network (FuNN) 15
 fuzzy rules 282
 fuzzy set 274
 fuzzy system modelling 274

G

gain token set (GTS) 62
 Gaussian distribution 127
 gbest model 275
 general public license (GPL) 285

general regression neural networks 256
 genetic algorithm (GA) 164, 186, 234, 243, 275
 genetic algorithm techniques 234
 grey-level histogram moments (GLHM) 259
 ground truth analysis 150
 group composition 168

H

Haiku 234, 236, 242, 247
 heuristic 217, 221
 hidden Markov models (HMMs) 76
 hierarchal networks 256
 hough circle detection 130
 hough transform 132
 human computer interface (HCI) 108
 human machine interfaces (HMI) 153

I

image pyramids 125
 impact boundaries 204
 impact detection 200
 impact gradient field (IGF) 220
 impact network 229
 impact sensing network 186, 217, 229
 inference rules 2
 informed maintenance grid 195
 infrared (IR) sensors 178
 input linguistic variables (ILV) 11
 intelligent agent 52
 intelligent decision support systems 1
 intelligent system 187, 188
 interaction 166
 interaction matrix (IM) 62

K

knowledge discovery 234, 236, 242
 knowledge management 2
 knowledge query and manipulation
 language (KQML) 53
 Kohonen networks 267

L

latent semantic kernel (LSK) 32, 40, 42
 lbest model 275

learning vector quantization (LVQNN) 267
 linear prediction coefficients (LPCs) 75
 linear predictive coding (LPC) 74
 literacy 100

M

Magna Science Adventure Centre 175
 maintenance programs 195
 Mamdani fuzzy model 3, 15, 20, 227, 282
 mammograms 251
 Mammographic Image Analysis Society
 (MIAS) 255
 manual classification 32
 matrix equation methods 61
 mean square error (MSE) 282
 Medicare benefit schedule (MBS) 30, 42
 medio-lateral oblique (MLO) 255
 membership function (MF) 11, 280, 293
 millibots 169
 minimum spanning tree (MST) 217
 minor damage 192
 momentum 258
 multi-agent interaction 166
 multi-agent system (MAS) 52, 53, 168, 189
 multi-input, multi-output (MIMO) 280
 multi-input, single-output (MISO) 280
 multi-layer perceptron (MLP) 78
 multi-sets of tokens 57

N

neuro-fuzzy (NF) 9
 Newtonian space 238
 Nickel-Cadmium battery 277, 284
 Nijmegen database 255
 non-critical damage 192
 non-understandable type set (NTS) 61

O

obstacle classification 136
 obstacle detection 125, 136, 148
 obstacle distillation 127
 obstacle tracking 128
 OIL 55
 online auction system 56

ontology 54
 ontology facilitators 56
 ontology language 55
 open environment 52
 optical flow 125
 orthogonal array (OA) 287

P

panoramic imaging 144
 particle swarm optimisation (PSO)
 273, 275, 278
 PasswordCourier® 83
 “pause” heuristic 221
 pedestrian detection 134
 pedestrian protection systems (PPS) 134
 pedestrian tracking 138
 Petri nets (PNs) 54, 57
 pheromone 159, 195
 Phraselator® 84
 place type (PT) 60
 place type set (PTS) 61
 predator-and-prey robots 175
 probabilistic neural network (PNN) 267
 PSO fuzzy modeler (MatLab-TM) 285
 PSO toolbox (MatLab-TM) 285

R

radial basis function neural network
 (RBFNN) 256, 267
 range finding 146
 reactivity 53
 reasoning 2
 receiver operating characteristic (ROC) 256
 recharging 171
 recurrent neural network 81
 regression trees 1
 required token set (RTS) 62
 response matrix 191, 229
 road sign detection 129
 robotics 157, 160
 rule firing strength (RFS) 12

S

safety-critical environment 191
 scaffolding 206

scaled conjugate gradient algorithm
 (SCGA) 1, 7
 scene understanding 130
 search and rescue 174
 segmentation 72, 123
 self organizing maps (SOMs) 79, 257, 267
 self-repair 214
 sensor 146, 187
 shape-replication 214
 signal-to-noise-ratio (SNR) 287
 simulated annealing optimization 259
 Singleton fuzzy model 286
 skin 196
 Smart Cars project 123
 smart skin 229
 Sobel edge detector 132
 soft computing (SC) 1, 6, 83
 spatial grey-level dependence (SGLD) 259
 speech processing 71
 speech production 69
 speech recognition 70, 102, 110
 Speechalator 84
 SpeechSecure™ 83
 speed sign detection 129
 stereo panoramic imaging 145
 stigmergy 166, 194
 structural health management (SHM) 186
 sub-critical damage 192, 202
 subsumption architecture 162
 sum-of-absolute-differences (SAD) 136
 supply chain formation (SCF) 64
 support vector machines (SVM) 33, 37,
 46, 136
 swarm intelligence (SI) 158, 275
 swarm robotics 159

T

tactical air combat decision support system
 (TACDSS) 2, 17
 tactical air combat environment (TACE) 1
 Tagaki-Sugeno fuzzy model 286
 Taguchi method 286, 292
 telecommunications 247
 template 133
 time delay neural network (TDNN) 81
 token 57

token colours 57, 60
top-down/bottom-up (TDBU) 194, 229
traffic control 171
travelling salesperson problem (TSP) 159,
217

U

uncertainty 2
understandable type set (UTS) 61
University of California, Irvine Machine
Learning Repository (UCI-ML)
245, 255

V

v-disparity algorithm 136
visualisation 237, 240, 242

Y

Yamabico robots 162

Z

zone of proximal development (ZPD) 88



Experience the latest full-text research in the fields
of Information Science, Technology & Management

InfoSci-Online

InfoSci-Online is available to libraries to help keep students, faculty and researchers up-to-date with the latest research in the ever-growing field of information science, technology, and management.

The InfoSci-Online collection includes:

- Scholarly and scientific book chapters
- Peer-reviewed journal articles
- Comprehensive teaching cases
- Conference proceeding papers
- All entries have abstracts and citation information
- The full text of every entry is downloadable in .pdf format

Some topics covered:

- Business Management
- Computer Science
- Education Technologies
- Electronic Commerce
- Environmental IS
- Healthcare Information Systems
- Information Systems
- Library Science
- Multimedia Information Systems
- Public Information Systems
- Social Science and Technologies

**InfoSci-Online
features:**

- Easy-to-use
- 6,000+ full-text entries
- Aggregated
- Multi-user access

"...The theoretical bent of many of the titles covered, and the ease of adding chapters to reading lists, makes it particularly good for institutions with strong information science curricula."

*— Issues in Science and
Technology Librarianship*



To receive your free 30-day trial access subscription contact:

Andrew Bundy

Email: abundy@idea-group.com • Phone: 717/533-8845 x29

Web Address: www.infosci-online.com

InfoSci-Online

Full Text • Cutting Edge • Easy Access

A PRODUCT OF  **IDEA GROUP INC.**

Publishers of Idea Group Publishing, Information Science Publishing, CyberTech Publishing, and IRM Press

infosci-online.com

Single Journal Articles and Case Studies Are Now Right at Your Fingertips!

Purchase any single journal article or
teaching case for only \$18.00!

Idea Group Publishing offers an extensive collection of research articles and teaching cases that are available for electronic purchase by visiting www.idea-group.com/articles. You will find over **980** journal articles and over **275** case studies from over 20 journals available for only \$18.00. The website also offers a new capability of searching journal articles and case studies by category. To take advantage of this new feature, please use the link above to search within these available categories:

- ◆ Business Process Reengineering
- ◆ Distance Learning
- ◆ Emerging and Innovative Technologies
- ◆ Healthcare
- ◆ Information Resource Management
- ◆ IS/IT Planning
- ◆ IT Management
- ◆ Organization Politics and Culture
- ◆ Systems Planning
- ◆ Telecommunication and Networking
- ◆ Client Server Technology
- ◆ Data and Database Management
- ◆ E-commerce
- ◆ End User Computing
- ◆ Human Side of IT
- ◆ Internet-Based Technologies
- ◆ IT Education
- ◆ Knowledge Management
- ◆ Software Engineering Tools
- ◆ Decision Support Systems
- ◆ Virtual Offices
- ◆ Strategic Information Systems Design, Implementation

You can now view the table of contents for each journal so it is easier to locate and purchase one specific article from the journal of your choice.

Case studies are also available through XanEdu, to start building your perfect coursepack, please visit www.xanedu.com.

For more information, contact cust@idea-group.com or 717-533-8845 ext. 10.

———— **www.idea-group.com** ————

 **IDEA GROUP INC.**

Idea Group REFERENCE

Authoritative | Innovative | Comprehensive

The Premier Reference Source for Information Science and Technology Research

Free online access for the life of edition with purchase of print copy

ENCYCLOPEDIA OF Gender and Information Technology




Edited by: Eileen Trauth
Two-Volume Set
March 2006; 1000+ pp
ISBN: 1-59140-815-6; US \$525.00 h/c
Pre-Publication Price: US \$445.00*
Online Electronic Access Only**:
US \$420.00

♦ The only reference work that provides an all-encompassing perspective on the way gender and information technology impact each other

♦ Includes over 200 entries with extensive reference sections listing additional works on gender and information technology in order to stimulate further research

ENCYCLOPEDIA OF Knowledge Management



Edited by: David Schwartz
September 2005; 902 pp
ISBN: 1-59140-573-4;
US \$275.00 h/c
Online Electronic Access Only**:
US \$220.00

"This volume is an essential reference source for ideas as to what needs to be addressed and what we have learned about knowledge management over the past few decades."

— Professor Laurence Prusak,
Babson College, USA

ENCYCLOPEDIA OF Multimedia Technology and Networking



Edited by: Margherita Pagani
Two-Volume Set
May 2005; 1218 pp
ISBN: 1-59140-561-0;
US \$495.00 h/c
Online Electronic Access Only**:
US \$396.00

"A remarkable collection of articles that are the finest in the field of multimedia technology and networking. This encyclopedia should be included in any library."

— Professor Leslie Leong,
Central Connecticut State University, USA



ENCYCLOPEDIA OF Distance Learning

Four-Volume Set
April 2005; 2418 pp
ISBN: 1-59140-555-6;
US \$995.00 h/c
Online Electronic Access Only**:
US \$796.00

"The four-volume set of the Encyclopedia of Distance Learning is full of quality work, knowledgeable contributions, and stimulation of ideas for research that is second to none."

— Professor Yair Levy,
Nova Southeastern University, USA,
in *International Journal of Distance Education Technologies*,
Volume 3, Issue 4

ENCYCLOPEDIA OF Information Science and Technology

Edited by: Mehdi Khosrow-Pour



Five-Volume Set
January 2005; 3807 pp
ISBN: 1-59140-553-X; US \$1125.00 h/c
Online Electronic Access Only**:
US \$900.00

"This set is best suited to academic libraries, particularly information systems and business collections, and library and information science collections. Special libraries and large public libraries may also find it of interest."

—Library Journal, May 15, 2005

Available
in print
and online.
Excellent additions
to your
institution's
library.

www.idea-group-ref.com

*Pre-pub price is good through one month after the publication date.

**Electronic access is for libraries only, and is good for the life of the edition